INDEX

Numerics

A

B

C

D

E

F

G

H

I

K

L

M

N

O

P

R

S

T

U

V

W

Preface

This preface describes who should read the *Cisco Wide Area Application Services Configuration Guide*, how it is organized, and its document conventions. It contains the following sections:

Audience

This guide is for experienced network administrators who are responsible for configuring and maintaining the Cisco Wide Area Application Services (WAAS) network.

You should be familiar with the basic concepts and terminology used in internetworking, and understand your network topology and the protocols that the devices in your network can use. You should also have a working knowledge of the operating systems on which you are running your WAAS network, such as Microsoft Windows, Linux, or Solaris.

Document Organization

This guide is organized as follows:

Chapter	Title	Description
Chapter 1	Introduction to Cisco WAAS	Provides an overview of the WAAS product and its features.
Chapter 2	Planning Your WAAS Network	Provides general guidelines and preparation information you should read before installing the WAAS product in your network.
Chapter 3	Using Device Groups and Device Locations	Describes how to create groups that make it easier to manage and configure multiple devices at the same time This chapter also covers device locations.

Chapter	Title	Description
Chapter 4	Configuring AppNav	Describes how to configure your WAAS network using the AppNav deployment model.
Chapter 5	Configuring Traffic Interception	Describes the WAAS software support for intercepting all TCP traffic in an IP-based network.
Chapter 6	Configuring Network Settings	Describes how to configure interfaces and basic network settings like DNS and CDP.
Chapter 7	Configuring Administrative Login Authentication, Authorization, and Accounting	Describes how to centrally configure administrative login authentication, authorization, and accounting for WAEs in your WAAS network.
Chapter 8	Creating and Managing Administrator User Accounts and Groups	Describes how to create device-based CLI accounts and roles-based accounts from the WAAS Central Manager GUI.
Chapter 9	Creating and Managing IP Access Control Lists for WAAS Devices	Describes how to centrally create and manage Internet Protocol (IP) access control lists (ACLs) for your WAEs.
Chapter 10	Configuring Other System Settings	Describes how to perform various other system configuration tasks such as specifying an NTP server and setting the time zone on a device.
Chapter 11	Using the WAE Device Manager GUI	Describes how to use the WAE Device Manager GUI to configure and manage individual WAEs in your network.
Chapter 12	Configuring File Services	Describes how to configure Common Internet File System (CIFS) acceleration, which allows branch office users to more efficiently access data stored at centralized data centers.
Chapter 13	Configuring Application Acceleration	Describes how to configure the application policies on your WAAS system that determine the types of application traffic that is accelerated over your WAN.
Chapter 14	Configuring Virtual Blades	Describes how to configure virtual blades, which emulate another computer in your WAAS device.
Chapter 15	Configuring the Network Analysis Module	Describes how to configure and use the Cisco Network Analysis Module (NAM) in the WAAS Central Manager.
Chapter 16	Maintaining Your WAAS System	Describes the tasks you may need to perform to maintain your WAAS system.
Chapter 17	Monitoring and Troubleshooting Your WAAS Network	Describes the monitoring and troubleshooting tools available in the WAAS Central Manager GUI that can help you identify and resolve issues with your WAAS system.

Chapter	Title	Description
Chapter 18	Configuring SNMP Monitoring	Describes how to configure SNMP traps, recipients, community strings and group associations, user security model groups, and user access permissions.
Appendix A	Predefined Optimization Policy	Lists the predefined applications and classifiers that WAAS will either optimize or pass through based on the policies that are provided with the system.
Appendix B	Transaction Log Format	Describes the transaction log format.

Document Conventions

Command descriptions use these conventions:

boldface font	Commands and keywords are in boldface.
italic font	Arguments for which you supply values are in italics.
[ ]	Elements in square brackets are optional.
[x | y | z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.

Screen examples use these conventions:

`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
`< >`	Nonprinting characters, such as passwords, are in angle brackets.
`[ ]`	Default responses to system prompts are in square brackets.
`!, #`	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:

Note Means reader *take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Tip Means t*he following information will help you solve a problem*. Tips might not be troubleshooting or even an action, but could help you save time.

Related Documentation

For additional information on the Cisco WAAS software and hardware, see the following documentation:

- *Release Note for Cisco Wide Area Application Services*
- *Cisco Wide Area Application Services Upgrade Guide*
- *Cisco Wide Area Application Services Command Reference*
- *Cisco Wide Area Application Services Quick Configuration Guide*
- *Cisco Wide Area Application Services Configuration Guide* (this manual)
- *Cisco Wide Area Application Services API Reference*
- *Cisco WAAS Troubleshooting Guide for Release 4.1.3 and Later*
- *Cisco Wide Area Application Services Monitoring Guide*
- *Cisco Wide Area Application Services vWAAS Installation and Configuration Guide*
- *Cisco WAAS Installation and Configuration Guide for Windows on a Virtual Blade*
- *Configuring WAAS Express*
- *Cisco WAAS on Service Modules for Cisco Access Routers*
- *Cisco SRE Service Module Configuration and Installation Guide*
- *Configuring Cisco WAAS Network Modules for Cisco Access Routers*
- *WAAS Enhanced Network Modules*
- *Regulatory Compliance and Safety Information for the Cisco Wide Area Virtualization Engines*
- *Cisco Wide Area Virtualization Engine 294 Hardware Installation Guide*
- *Cisco Wide Area Virtualization Engine 594 and 694 Hardware Installation Guide*
- *Cisco Wide Area Virtualization Engine 7541, 7571, and 8541 Hardware Installation Guide*
- *Cisco Wide Area Virtualization Engine 274 and 474 Hardware Installation Guide*
- *Cisco Wide Area Virtualization Engine 574 Hardware Installation Guide*
- *Regulatory Compliance and Safety Information for the Cisco Content Networking Product Series*
- *Cisco Wide Area Application Engine 7341, 7371, and 674 Hardware Installation Guide*
- *Installing the Cisco WAE Inline Network Adapter*

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

CHAPTER 1

Introduction to Cisco WAAS

This chapter provides an overview of the Cisco WAAS solution and describes the main features that enable WAAS to overcome the most common challenges in transporting data over a wide area network.

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE and WAVE appliances, WAE Network Modules (the NME-WAE family of devices), SM-SRE modules running WAAS, and vWAAS instances.

This chapter contains the following sections:

About Cisco WAAS

The WAAS system consists of a set of devices called wide area application engines (WAEs) that work together to optimize TCP traffic over your network. When client and server applications attempt to communicate with each other, the network intercepts and redirects this traffic to the WAEs so that they can act on behalf of the client application and the destination server. The WAEs examine the traffic and use built-in optimization policies to determine whether to optimize the traffic or allow it to pass through your network unoptimized.

WAAS version 5.0 introduces a new AppNav deployment model that greatly reduces dependency on the intercepting switch or router by taking the responsibility of distributing traffic among WAAS devices for optimization. WAAS appliances with AppNav Controller Interface Modules operate in a special AppNav Controller mode with AppNav policies controlling traffic flow to WAAS devices doing optimization. The AppNav model is well suited to data center deployments and addresses many of the challenges of WAN optimization in this environment.

You can deploy WAAS in the new AppNav model or in the traditional model without using AppNav Controllers.

You use the WAAS Central Manager GUI to centrally configure and monitor the WAEs and optimization policies in your network. You can also use the WAAS Central Manager GUI to create new optimization policy rules so that the WAAS system can optimize custom applications and less common applications.

Cisco WAAS helps enterprises meet the following objectives:

- Provide branch office employees with LAN-like access to information and applications across a geographically distributed network.
- Migrate application and file servers from branch offices into centrally managed data centers.
- Minimize unnecessary WAN bandwidth consumption through the use of advanced compression algorithms.
- Virtualize print and other local services to branch office users. Cisco WAAS allows you to configure a WAE with Windows in a virtual blade so that you do not need to deploy a dedicated system to provide local services such as Print Services, Active Directory Services, DNS, and DHCP services.
- Improve application performance over the WAN by addressing the following common issues:
  - Low data rates (constrained bandwidth)
  - Slow delivery of frames (high network latency)
  - Higher rates of packet loss (low reliability)

Note A WAAS Express device, which is a Cisco router with WAAS Express functionality enabled, can interoperate with WAE devices. A WAAS Express device provides basic WAN optimization and some application optimization but no virtualization. For more information on WAAS Express, see *Configuring WAAS Express*.

A virtual WAAS (vWAAS) instance is a virtual WAAS appliance running on a VMware virtual machine and providing all of the same features as a WAAS appliance. A WAAS Central Manager can manage WAEs, WAAS Express devices, and vWAAS instances all in the same WAAS network. For more information on vWAAS, see the *Cisco Wide Area Application Services vWAAS Installation and Configuration Guide*.

This section contains the following topics:

Cisco WAAS Overcomes Common WAN Challenges

Table 1-1 describes how Cisco WAAS uses a combination of TCP optimization techniques and application acceleration features to overcome the most common challenges associated with transporting traffic over a WAN.

Table 1-1 Cisco WAAS Solution

WAN Issue	WAAS Solution
High network latency	Intelligent protocol adapters reduce the number of roundtrip responses common with chatty application protocols.
Constrained bandwidth	Data caching provided with the file services feature and data compression reduce the amount of data sent over the WAN, which increases data transfer rates. These solutions improve application response time on congested links by reducing the amount of data sent across the WAN.

Table 1-1 Cisco WAAS Solution (continued)

WAN Issue	WAAS Solution
Poor link utilization	TCP optimization features improve network throughput by reducing the number of TCP errors sent over the WAN and maximizing the TCP window size that determines the amount of data that a client can receive at one time.
Packet loss	Optimized TCP stack in WAAS overcomes the issues associated with high packet loss and protects communicating end points from the state of the WAN.

Traffic Optimization Process

Figure 1-1 shows the process that Cisco WAAS follows to optimize application traffic.

Figure 1-1 Traffic Optimization Process

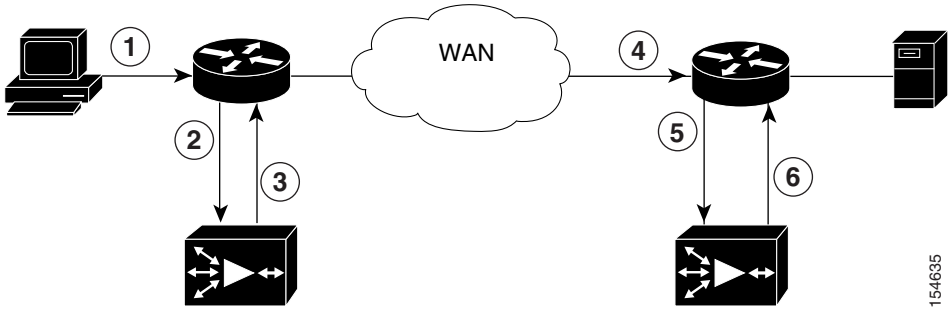

The following steps describe how your WAAS network optimizes a connection between a branch office client and a destination server:

1. A branch office client attempts to connect to the destination server over the native application port.
2. The WAAS network uses WCCP or PBR to intercept the client request, or if deployed on an inline WAE, WAAS can intercept the request directly using inline mode. For more information on inline mode, see the "Using Inline Mode Interception" section on page 5-42.
3. The branch WAE performs the following actions:

- Examines the parameters in the traffic's TCP headers and then refers to the optimization policies to determine if the intercepted traffic should be optimized. Information in the TCP header, such as the source and destination IP address and port, allows the branch WAE to match the traffic to an optimization policy rule. For a list of predefined policy rules, see Appendix A, "Predefined Optimization Policy."
- If the branch WAE determines that the traffic should be optimized, it adds information to the TCP header that informs the next WAE in the network path to optimize the traffic.

4. The branch WAE passes along the client request through the network to its original destination server.
5. The data center WAE performs the following actions:

- Intercepts the traffic going to the destination server.

- Establishes an optimized connection with the branch WAE. If the data center WAE has optimization disabled, then an optimized connection will not be established and the traffic passes over the network unoptimized.

  In an AppNav deployment, an AppNav Controller intercepts the traffic in the data center and distributes it to a WAAS node that establishes an optimized connection with the branch WAE. For more information on an AppNav deployment, see Chapter 4, "Configuring AppNav."

6. WAAS optimizes subsequent traffic between the branch WAE and data center WAE for this connection.

Cisco WAAS does not optimize traffic in the following situations:

- The WAE intercepts non-TCP traffic (such as UDP or ICMP).
- The WAE is overloaded and does not have the resources to optimize the traffic.
- The intercepted traffic matches an optimization or AppNav policy rule that specifies to pass the traffic through unoptimized.

Note If unoptimized traffic reaches a WAE, the WAE forwards the traffic in pass-through mode without affecting the performance of the application using the passed-through connection.

Key Services of Cisco WAAS

Cisco WAAS contains the following services that help optimize traffic over your wide area network:

- TFO Optimization, page 1-4
- Compression, page 1-6
- Application-Specific Acceleration, page 1-6
- File Services for Desktop Applications, page 1-8
- WAAS Print Services, page 1-9
- Virtualization, page 1-10

Note WAAS Express devices provide basic optimization and compression services and some application acceleration.

TFO Optimization

Cisco WAAS uses a variety of transport flow optimization (TFO) features to optimize TCP traffic intercepted by the WAAS devices. TFO protects communicating clients and servers from negative WAN conditions, such as bandwidth constraints, packet loss, congestion, and retransmission.

TFO includes the following optimization features:

- Windows Scaling, page 1-5
- TCP Initial Window Size Maximization, page 1-5
- Increased Buffering, page 1-5
- Selective Acknowledgment, page 1-5

- BIC TCP, page 1-6

Windows Scaling

Windows scaling allows the receiver of a TCP packet to advertise that its TCP receive window can exceed 64 KB. The receive window size determines the amount of space that the receiver has available for unacknowledged data. By default, TCP headers limit the receive window size to 64 KB, but Windows scaling allows the TCP header to specify receive windows of up to 1 GB.

Windows scaling allows TCP endpoints to take advantage of available bandwidth in your network and not be limited to the default window size specified in the TCP header.

For more information about Windows scaling, refer to RFC 1323.

TCP Initial Window Size Maximization

WAAS increases the upper bound limit for TCP's initial window from one or two segments to two to four segments (approximately 4 KB). Increasing TCP's initial window size provides the following advantages:

- When the initial TCP window is only one segment, a receiver that uses delayed ACKs is forced to wait for a timeout before generating an ACK response. With an initial window of at least two segments, the receiver generates an ACK response after the second data segment arrives, eliminating the wait on the timeout.
- For connections that transmit only a small amount of data, a larger initial window reduces the transmission time. For many e-mail (SMTP) and web page (HTTP) transfers that are less than 4 KB, the larger initial window reduces the data transfer time to a single round trip time (RTT).
- For connections that use large congestion windows, the larger initial window eliminates up to three RTTs and a delayed ACK timeout during the initial slow-start phase.

For more information about this optimization feature, see RFC 3390.

Increased Buffering

Cisco WAAS enhances the buffering algorithm used by the TCP kernel so that WAEs can more aggressively pull data from branch office clients and remote servers. This increased buffer helps the two WAEs participating in the connection keep the link between them full, increasing link utilization.

Selective Acknowledgment

Selective Acknowledgement (SACK) is an efficient packet loss recovery and retransmission feature that allows clients to recover from packet losses more quickly than the default recovery mechanism used by TCP.

By default, TCP uses a cumulative acknowledgement scheme that forces the sender to either wait for a roundtrip to learn if any packets were not received by the recipient or to unnecessarily retransmit segments that may have been correctly received.

SACK allows the receiver to inform the sender about all segments that have arrived successfully, so the sender only needs to retransmit the segments that have actually been lost.

For more information about SACK, see RFC 2018.

BIC TCP

Binary Increase Congestion (BIC) TCP is a congestion management protocol that allows your network to recover more quickly from packet loss events.

When your network experiences a packet loss event, BIC TCP reduces the receiver's window size and sets that reduced size as the new value for the minimum window. BIC TCP then sets the maximum window size value to the size of the window just before the packet loss event occurred. Because packet loss occurred at the maximum window size, the network can transfer traffic without dropping packets whose size falls within the minimum and maximum window size values.

If BIC TCP does not register a packet loss event at the updated maximum window size, that window size becomes the new minimum. If a packet loss event does occur, that window size becomes the new maximum. This process continues until BIC TCP determines the new optimum minimum and maximum window size values.

Compression

Cisco WAAS uses the following compression technologies to help reduce the size of data transmitted over your WAN:

- Data Redundancy Elimination (DRE)
- LZ compression

These compression technologies reduce the size of transmitted data by removing redundant information before sending the shortened data stream over the WAN. By reducing the amount of transferred data, WAAS compression can reduce network utilization and application response times.

When a WAE uses compression to optimize TCP traffic, it replaces repeated data in the stream with a much shorter reference, then sends the shortened data stream out across the WAN. The receiving WAE uses its local redundancy library to reconstruct the data stream before passing it along to the destination client or server.

The WAAS compression scheme is based on a shared cache architecture where each WAE involved in compression and decompression shares the same redundancy library. When the cache that stores the redundancy library on a WAE becomes full, WAAS uses a FIFO algorithm (first in, first out) to discard old data and make room for new.

LZ compression operates on smaller data streams and keeps limited compression history. DRE operates on significantly larger streams (typically tens to hundreds of bytes or more) and maintains a much larger compression history. Large chunks of redundant data is common in file system operations when files are incrementally changed from one version to another or when certain elements are common to many files, such as file headers and logos.

Application-Specific Acceleration

In addition to the TCP optimization features that speed the flow of traffic over a WAN, Cisco WAAS includes these application acceleration features:

- Operation prediction and batching—Allows a WAAS device to transform a command sequence into a shorter sequence over the WAN to reduce roundtrips.
- Intelligent message suppression—Decreases the response time of remote applications. Even though TFO optimizes traffic over a WAN, protocol messages between branch office clients and remote servers can still cause slow application response time. To resolve this issue, each WAAS device

contains application proxies that can respond to messages locally so that the client does not have to wait for a response from the remote server. The application proxies use a variety of techniques including caching, command batching, prediction, and resource prefetch to decrease the response time of remote applications.

- CIFS caching—Allows a WAAS device to reply to client requests using locally cached data instead of retrieving this data from remote file and application servers.
- Preposition—Allows a WAAS device to prefetch resource data and metadata in anticipation of a future client request. (Only the CIFS accelerator supports prepositioning.)

Cisco WAAS uses application-intelligent software modules to apply these acceleration features.

In a typical Common Internet File System (CIFS) application use case, the client sends a large number of synchronous requests that require the client to wait for a response before sending the next request. Compressing the data over the WAN is not sufficient for acceptable response time.

For example, when you open a 5 MB Word document, about 700 CIFS requests (550 read requests plus 150 other requests) are produced. If all these requests are sent over a 100 ms round-trip WAN, the response time is at least 70 seconds (700 x 0.1 seconds).

WAAS application acceleration minimizes the synchronous effect of the CIFS protocol, which reduces application response time. Each WAAS device uses optimization policies to match specific types of the traffic to an application and to determine whether that application traffic should be optimized and accelerated.

The following WAAS application accelerators are available:

- SMB—Accelerates CIFS traffic exchanged with a remote file server. Supports the SMB 1.0, 2.0, and 2.1 protocols for CIFS traffic and signed SMB traffic. For more information, see the "File Services for Desktop Applications" section on page 1-8.
- CIFS—Accelerates CIFS traffic exchanged with a remote file server. Supports the SMB 1.0 protocol for CIFS traffic. For more information, see the "File Services for Desktop Applications" section on page 1-8.

Note The SMB and CIFS application accelerators both handle CIFS traffic but have slightly different features. You must choose one or the other to operate on WAAS peer devices because they cannot operate simultaneously on the same device and both peers must use the same accelerator.

- NFS—Accelerates Network File System (NFS) version 3 traffic exchanged with a remote file server. Secure NFS traffic is not accelerated.
- ICA—Accelerates Independent Computing Architecture (ICA) traffic that is used to access a virtual desktop infrastructure (VDI).
- HTTP—Accelerates HTTP and HTTPS traffic.
- SSL—Accelerates encrypted Secure Sockets Layer (SSL) and Transport Layer Security (TLS) traffic. The SSL accelerator provides traffic encryption and decryption within WAAS to enable end-to-end traffic optimization. The SSL accelerator also provides secure management of the encryption certificates and keys.
- MAPI—Accelerates Microsoft Outlook Exchange traffic that uses the Messaging Application Programming Interface (MAPI) protocol. Microsoft Outlook 2000–2010 clients are supported. Secure connections that use message authentication (signing) or encryption are accelerated. MAPI over HTTP is not accelerated.

- Video—Accelerates Windows Media live video broadcasts that use RTSP over TCP. The video accelerator automatically splits one source video stream from the WAN into multiple streams to serve multiple clients on the LAN. The video accelerator automatically causes a client requesting a UDP stream to do a protocol rollover to use TCP (if both the client and server allow TCP).
- Windows Print—Accelerates print traffic between clients and a Windows print server located in the data center. Signed Server Message Block (SMB) traffic is optimized by transport level optimizations (TFO, DRE, and LZ). The Windows print accelerator supports Windows 2000, Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2 print servers. It supports clients running Windows 2000, Windows XP, Windows Vista, and Windows 7. The Windows Print accelerator operates only when the CIFS application accelerator is enabled.

Note WAAS Express devices provide application acceleration for CIFS/SMB, HTTP, and SSL traffic.

To enable or disable application accelerators, see the "Enabling and Disabling the Global Optimization Features" section on page 13-3.

You must enable the accelerator on both of the peer WAEs at either end of a WAN link for all application accelerators to operate.

File Services for Desktop Applications

The file services (SMB and CIFS accelerators) feature allows a WAE to more quickly fulfill a client's requests instead of sending every request over the WAN to the file server. By fulfilling the client's requests locally, the WAE minimizes the traffic sent over the WAN and reduces the time it takes branch office users to access files and many desktop applications, allowing enterprises to consolidate their important information into data centers.

For more information, see Chapter 12, "Configuring File Services."

Note Legacy mode WAFS is no longer supported. Legacy WAFS users must migrate to the SMB or CIFS accelerators.

This section contains the following topics:

File Services Features

File Services include the following features:

- Data coherency and concurrency—Ensures data integrity across the WAAS system by managing the freshness of the data (coherency) and controlling the access to the data by multiple clients (concurrency).
- Automatic discovery—Allows you to use file services without having to register individual file servers in the WAAS Central Manager. With the automatic discovery feature, the WAAS device will automatically discover and connect to a new file server when a CIFS request is received.

- Prepositioning—Allows system administrators to proactively "push" frequently used files from the central file server into the cache of selected WAEs, which provides users with faster first-time file access, and makes more efficient use of available bandwidth. Prepositioning is supported only by the CIFS application accelerator.

Role of the Edge WAE

The Edge WAE is a client-side, file-caching device that serves client requests at remote sites and branch offices. The device is deployed at each branch office or remote campus, replacing file and print servers and giving local clients fast, near-LAN read and write access to a cached view of the centralized storage. By caching the data most likely to be used at these sites, Edge WAEs greatly reduce the number of requests and the volume of data that must be transferred over the WAN between the data center and the edge.

When requests for data that is not located in the cache are received, the Edge WAE encapsulates the original CIFS request using a TCP/IP-based protocol, compresses it, and sends it over the WAN to the Core WAE. Data returned from the data center is distributed by the Edge WAE to the end user who requested it.

Role of the Core WAE

The Core WAE is a server-side component that resides at the data center and connects directly to one or more file servers or network-attached storage (NAS). Core WAEs are placed between the file servers at the data center and the WAN connecting the data center to the enterprise's remote sites and branch offices. Requests received from Edge WAEs over the WAN are translated by the Core WAE into its original file server protocol and forwarded to the appropriate file server. The data center Core WAEs can provide load balancing and failover support.

When the data is received from the file server, the Core WAE encapsulates and compresses it before sending it over the WAN back to the Edge WAE that requested it. Core WAEs can be arranged in logical clusters to provide scalability and automatic failover capabilities for high-availability environments.

WAAS Print Services

The WAAS software includes the following print services options:

- Windows print accelerator—Use this option when you have a print server in a data center and branch clients are printing to local or remote printers. This service accelerates print traffic between clients and a Windows print server located in the data center. This option requires no configuration but does require that both the CIFS application accelerator and Windows print acceleration be enabled. For more information, see the "Enabling and Disabling the Global Optimization Features" section on page 13-3.
- Virtual blade based print server—Use this option when you want to deploy a local print server in the branch office but without installing separate print server hardware. You can install a Windows print server in a virtual blade on the branch WAE, which allows you to manage printing by using standard Windows print server functionality. For more information, see Chapter 14, "Configuring Virtual Blades."

Note The legacy print services feature is no longer supported. Legacy print services users must migrate to another print services option.

These services eliminate the need for a separate hardware print server in the branch office. WAAS print services are available for Windows clients and work with any IP-based network printer.

Virtualization

The WAAS software allows you to configure a virtual blade, which allows you to add services running in their own operating environments to your WAAS system. For example, you could configure a virtual blade in a WAE device to run Windows services such as Print Services, Active Directory Services, DNS, and DHCP services.

A WAAS virtual blade provides an emulated hardware environment within your WAE device that acts as a generic computer. You can install an operating system and applications to work with your WAAS system and provide additional services for the users on your network. For more information, see Chapter 14, "Configuring Virtual Blades."

Overview of the WAAS Interfaces

The WAAS software provides the following interfaces to help you manage, configure, and monitor the various elements of your WAAS network:

- WAAS Central Manager GUI, page 1-10
- WAAS Central Manager Monitoring API, page 1-17
- WAE Device Manager GUI, page 1-17
- WAAS CLI, page 1-18
- WAAS CLI, page 1-18

WAAS Central Manager GUI

Every WAAS network must have one primary WAAS Central Manager device that is responsible for managing the other WAAS devices in your network. The WAAS Central Manager devices hosts the WAAS Central Manager GUI, a Web-based interface that allows you to configure, manage, and monitor the WAAS devices in your network. The WAAS Central Manager resides on a dedicated WAE device.

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

- Configure system and network settings for an individual WAAS device, vWAAS device, WAAS Express device, device group, AppNav Controller, and AppNav Cluster.
- Create and edit optimization policies that determine the action that a WAAS device performs when it intercepts specific types of traffic.
- Create and edit AppNav policies that determine how AppNav Controllers distribute traffic to optimizing WAAS nodes.
- Configure file services and set up file preposition policies (preposition works only with the CIFS application accelerator).
- Create device groups that help you manage and configure multiple WAEs at the same time.
- View detailed reports about the optimized traffic in your WAAS network.

Note You cannot enable optimization and application acceleration services on a WAE that has been configured as a WAAS Central Manager. The purpose of the WAAS Central Manager is to configure, monitor, and manage the WAEs in your network.

This section contains the following topics:

Accessing the WAAS Central Manager GUI

To access the WAAS Central Manager GUI, enter the following URL in your web browser:

https://*WAE_Address*:8443/

The *WAE_Address* value is the IP address or hostname of the WAAS Central Manager device.

The default administrator username is *admin* and the password is *default*. For information on creating accounts and changing passwords, see Chapter 8, "Creating and Managing Administrator User Accounts and Groups."

Ensure that your web browser is set to use Unicode (UTF-8) character encoding.

Note When using Internet Explorer to access the Central Manager GUI, you may see a "Choose a digital certificate" dialog. Click **Cancel** to proceed to the Central Manager login screen.

You may also see a browser security warning that there is a problem with the website's security certificate. This happens because the Central Manager uses a self-signed certificate. Click on the link **Continue to this website (not recommended)**. You can permanently install the certificate to avoid this error in the future. To install the certificate in Internet Explorer 8, click the red **Certificate Error** button in the address bar and choose **View Certificates**. Click **Install Certificate**, then click **Next**. Select Automatically select the certificate store based on the type of certificate and click **Next**, click **Finish**, then click **Yes** on the security warning, click **OK** on the acknowledgement, and click **OK** on the Certificate dialog. The certificate installation procedure differs depending on the browser.

If you are using Internet Explorer to access the Central Manager GUI, we strongly recommend that you install the Google Chrome Frame plug-in to provide better performance. When you log into the Central Manager the first time, you are prompted to install Google Chrome Frame. Choose a language, click **Get Google Chrome Frame**, and follow the prompts to download and install the plug-in. If you do not want to install the plugin, click the link to continue without installing Google Chrome Frame.

You can configure the WAAS Central Manager GUI to limit the number of concurrent sessions permitted for a user. The number of concurrent sessions is unlimited by default. To change the number of permitted concurrent sessions, set the System.security.maxSimultaneousLogins property, as described in the "Modifying the Default System Configuration Properties" section on page 10-17.

Note A user must log off the Central Manager to end a session. If a user closes the browser or connection without logging off, the session is not closed until after it times out (in 10 minutes by default, up to a possible maximum of 120 minutes). If the number of concurrent sessions permitted also is exceeded for that user, there is no way for that user to regain access to the Central Manager GUI until after the timeout expires.

Note After an upgrade, downgrade, or new installation, you must first clear the cache in your browser, close the browser, and restart the browser session to the WAAS Central Manager.

Components of the WAAS Central Manager GUI

Figure 1-2 shows the main components of the WAAS Central Manager GUI.

Figure 1-2 Components of the WAAS Central Manager GUI

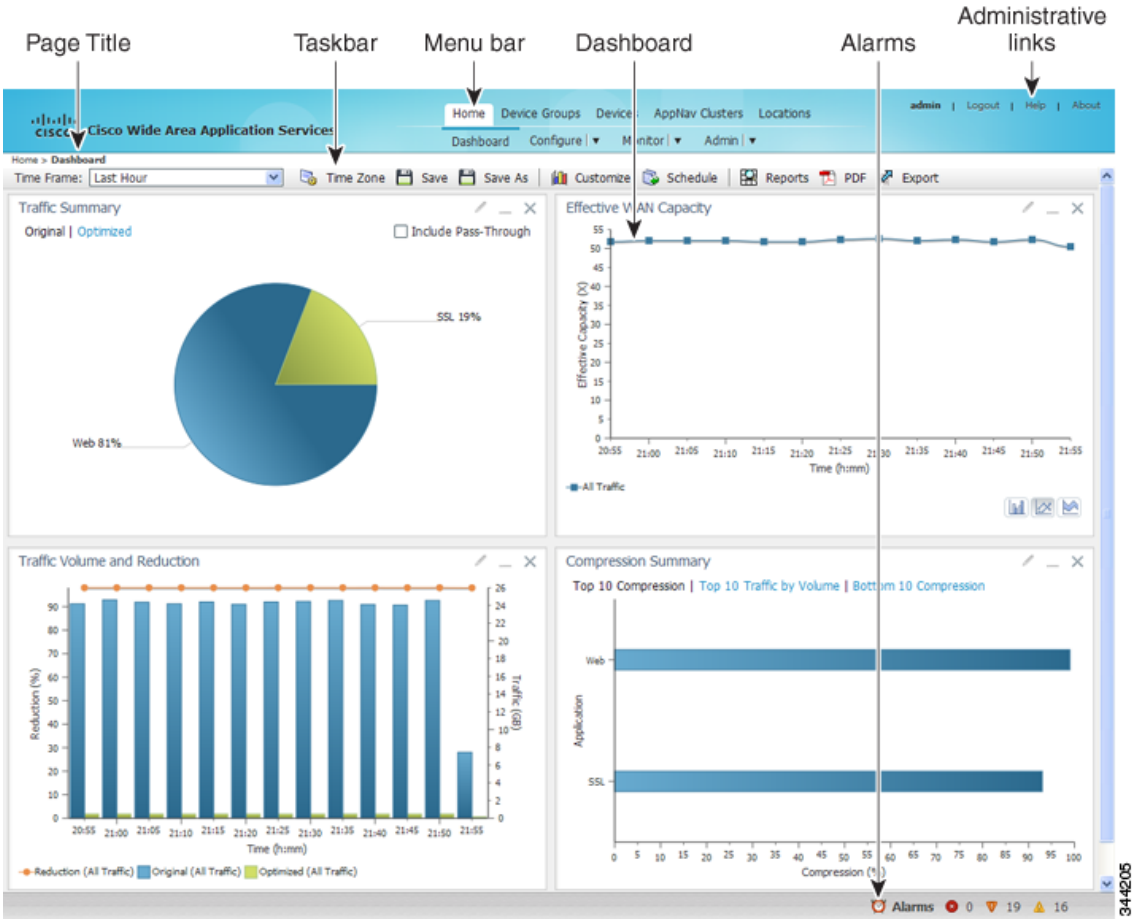

The WAAS Central Manager GUI includes the following main components:

- Page title—Displays the title of the page being viewed and breadcrumb links to ease navigation back to previous levels in the hierarchy. (Breadcrumb links are shown in Figure 1-3.)

- Menu bar—The top level contains menus that allow you to choose the context. The lower level contains menus that group the WAAS Central Manager functions available within the chosen context. For more information, see the "WAAS Central Manager Menus" section on page 1-14.
- Taskbar—Contains labeled icons that perform various functions depending on the content shown in the dashboard. For more information, see the "WAAS Central Manager Taskbar Icons" section on page 1-15.
- Dashboard—Displays the main content, which changes depending on the function that is chosen in the menu.
- Administrative links—Includes these navigation links:
  - Logout—Logs out the current user from the WAAS Central Manager.
  - Help—Opens a separate window with the WAAS context sensitive help.
  - About—Displays the WAAS About screen that shows the Central Manager version number.
- Alarms—Opens the alarm panel, which displays alarms in your WAAS network.

The top level of the menu bar allows you to choose one of the five contexts available in the WAAS Central Manager GUI:

- Home—Click to go to the global context, with no particular device group, device, AppNav Cluster, or location chosen.
- Device Groups—Choose a device group from this menu to enter the device group context. The page title and the first menu on the lower level displays the name of the chosen device group.
- Devices—Choose a device from this menu to enter the device context. The page title and the first menu on the lower level displays the name of the chosen device, as shown in Figure 1-3.
- AppNav Clusters—Choose an AppNav Cluster from this menu to enter the AppNav Cluster context. The page title and the first menu on the lower level displays the name of the chosen AppNav Cluster.
- Locations—Choose a location from this menu to enter the location context. The page title and the first menu on the lower level displays the name of the chosen location.

Figure 1-3 WAAS Central Manager Device Context

The WAAS Central Manager GUI includes the following items to help you navigate:

- Breadcrumbs to current location—Displays the path to your current location in the menu structure. You can click the Devices link to return to the All Devices page. If you are in the device group context, this link is named Device Groups and it returns you to the All Device Groups page. If you are in the AppNav Cluster context, this link is named AppNav Clusters and it returns you to the All AppNav Clusters page. If you are in the location context, this link is named Locations and it returns you to the All Locations page.

- Entity name—The first menu in the lower level of the menu bar shows the name of the chosen device group, device, AppNav Cluster, or location.
- Context menus—The top level of the menu bar contains menus that allow you to switch easily to any entity in any context. You can search for an item by entering part of its name in the search box at the top and clicking the magnifying glass icon or pressing **Enter**. The list is filtered to include only entities that contain the search string. The top entry in each menu is All *Entities*, which takes you to a full window that lists all entities of the selected type, has more advanced search functions, and has taskbar icons that perform functions appropriate to the entity group. You can also click the context menu name to go to the listing window.

  In the Devices and AppNav Clusters menus, a small target icon appears when you hover over a device or cluster name. Place your cursor over the target icon to open a pop-up that shows the device or cluster status (see Figure 1-4).

Figure 1-4 Devices Context Menu

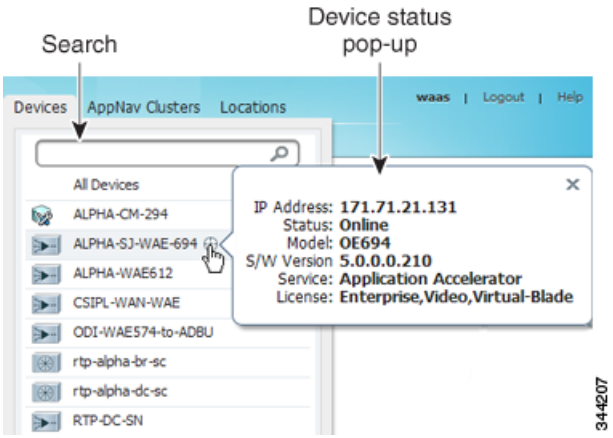

WAAS Central Manager Menus

The WAAS Central Manager menu bar contains two levels of menus:

- Top level—Contains menus that allow you to switch to any entity in any context.
- Lower level—Contains menus that group the WAAS Central Manager functions available within the chosen context. Table 1-2 describes the menus in the lower menu bar.

Menus contain different functions when a particular device, device group, AppNav Cluster, or location is selected than when you are in the global context.

Some menu options contain submenus. Hover over the triangle to the right of the menu option name to open the submenu.

Note The functions available for WAAS Express devices are a subset of those available for other WAAS devices; some functions are not available on WAAS Express devices.

Table 1-2 Menu Descriptions

Menu	Description
Dashboard or *Device, Device group, AppNav Cluster, or Location name*	In the global context, allows you to go to the dashboard for your WAAS network. In a context other than global, this menu is named with the entity name and allows you to activate devices, view users, assign groups or devices, or view the dashboard or home screen of the entity.
Configure	Allows you to configure WAAS services and settings.
Monitor	Allows you to see network traffic and other charts and reports to monitor the health and performance of your WAAS network. Allows you to manage and schedule reports for your WAAS network. Contains troubleshooting tools.
Admin	Allows you to manage user accounts, passwords, secure store, licenses, and virtual blades, update the WAAS software, and view system logs and messages.

WAAS Central Manager Taskbar Icons

Table 1-3 describes the taskbar icons in the WAAS Central Manager GUI.

Table 1-3 Taskbar Icon Descriptions (continued)

Taskbar Icon	Function
Common icons	
(Refresh)	Refreshes the current page of the WAAS Central Manager GUI.
(Delete)	Deletes a WAAS element, such as a device, and device group.
(Create or Add)	Creates a new WAAS element such as a report.
(Edit)	Edits a WAAS element such as interface settings.
(Advanced Search)	Filters the information in a table to make it easier to locate a specific item.
(View All)	Displays all items in a table on a single page instead of displaying those items over multiple pages.
(Print or Print Table)	Prints the information.
(PDF)	Creates a PDF of the information.
(Assign All)	Selects all valid items in a table. For example, if you are distributing print drivers to a WAAS print server, you can click this icon to select all drivers in the list that the print server should download.
(Remove All)	Deselects all selected items in a table.

Table 1-3 Taskbar Icon Descriptions (continued)

Taskbar Icon	Function
Devices and Device Group Icons	
(Activate All Inactive Devices)	Activates all the inactive WAAS and WAAS Express devices in your WAAS network. For more information, see the "Activating All Inactive WAAS Devices" section on page 16-34.
(Force Update, Request FullUpdate)	Reapplies the device configuration as seen in the WAAS Central Manager GUI to the device. Normally, changes made in the WAAS Central Manager GUI are applied to the device as soon as the configuration is submitted. From time to time, however, a CLI error or some other error on the device can cause the configuration on the device to differ from what is seen in the WAAS Central Manager GUI. The Force Full Database Update icon applies the full configuration that the WAAS Central Manager has for the device to be updated to the device and the configuration reapplied. When using the Request FullUpdate icon from the device group window, the full device configuration is reapplied to each device in the device group. Group settings do not overwrite device-specific settings. You can view device CLI errors in the System Message window described in the "Viewing the System Message Log" section on page 17-56. The Force Full Database Update icon appears on the Device Dashboard window, described in the "Device Dashboard Window" section on page 17-8. The Request FullUpdate icon appears on the Modifying Device Group window. These functions do not apply to WAAS Express devices.
(Reload)	Reboots a WAE or device group depending on the location in the WAAS Central Manager GUI. For more information, see the "Rebooting a Device or Device Group" section on page 16-35. Reload is not available for WAAS Express devices.
(Force Group Settings)	Forces the device group configuration across all devices in that group. For more information, see the "Forcing Device Group Settings on All Devices in the Group" section on page 3-7.
(Apply Defaults)	Applies the default settings to the fields on the window.
(Export Table)	Exports table information into a CSV file.
(Override Group Settings)	Allows you to specify device-specific settings that override the group settings for the device. For more information, see the "Overriding the Device Group Settings on a Device" section on page 3-8.
(Deactivate Device)	Deactivates a WAAS or WAAS Express device.
(Update Application Statistics)	Updates the application statistics.

Table 1-3 Taskbar Icon Descriptions (continued)

Taskbar Icon	Function
(Delete All)	Deletes all WAAS elements of a particular type, such as IP ACL conditions.
(Display All Devices)	Displays all WAE devices or device groups.
(Configure Dashboard Display)	Allows you choose which charts to display in the Device Dashboard window.
(Copy Settings)	Copies interception settings to other devices (not available for inline interception).
Acceleration Icons	
(Restore Default Policies and Classifiers)	Restores the default predefined optimization policy rules on the device or device group. For more information, see the "Restoring Optimization Policies and Class Maps" section on page 13-58.
(View Topology)	Displays the topology map that shows all the TFO connections among your WAE devices. For more information, see the "Topology Report" section on page 17-40.
(Navigate to Application Configuration Page)	Displays the configuration page used to create applications. For more information, see the "Viewing a List of Applications" section on page 13-56.
System Message Log Icons	
(Truncate Table)	Allows you to truncate the system message log based on size, date, or message content. For more information, see the "Viewing the System Message Log" section on page 17-56.

WAAS Central Manager Monitoring API

The WAAS Central Manager monitoring application programming interface (API), provides a programmable interface for system developers to integrate with customized or third-party monitoring and management applications. The Central Manager monitoring API communicates with the WAAS Central Manager to retrieve status information and monitoring statistics.

The Central Manager monitoring API is a Web Service implementation. Web Service is defined by the W3C standard as a software system designed to support interoperable machine-to-machine (client and server) interaction over the network. The client and server communication follows the Simple Object Access Protocol or Service Oriented Architecture Protocol (SOAP) standard.

For more information on the monitoring API, see the *Cisco Wide Area Application Services API Reference*.

WAE Device Manager GUI

The WAE Device Manager is a web-based management interface that allows you to configure, manage, and monitor an individual WAE device in your network. In some cases, the same device settings exist in both the WAE Device Manager and the WAAS Central Manager GUI. For this reason, we recommend that you always configure device settings from the WAAS Central Manager GUI when possible.

In some situations, you might need to use the WAE Device Manager GUI to perform certain tasks. For example, starting, stopping, and restarting the CIFS accelerator service can only be performed from the WAE Device Manager GUI and not from the WAAS Central Manager GUI.

For more information about the tasks you can perform from the WAE Manager, see Chapter 11, "Using the WAE Device Manager GUI."

To access the WAE Device Manager for a specific device, go to the following URL:

https://*Device IP Address*:8443/mgr

Figure 1-5 shows an example of the WAE Device Manager window.

Figure 1-5 WAE Device Manager Window

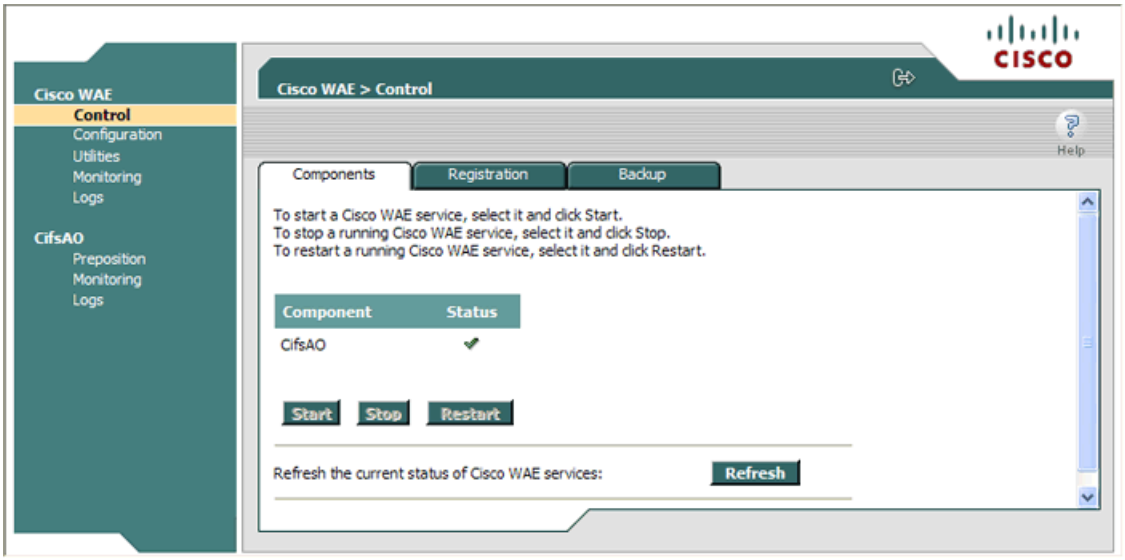

WAAS CLI

The WAAS CLI allows you to configure, manage, and monitor WAEs on a per-device basis through a console connection or a terminal emulation program. The WAAS CLI also allows you to configure certain features that are supported only through the CLI (for example, configuring the Lightweight Directory Access Protocol [LDAP] signing on a WAE). We strongly recommend that you use the WAAS Central Manager GUI instead of the WAAS CLI, whenever possible.

Note You must wait for approximately 10 minutes (two data feed poll cycles) after registering a WAE with the WAAS Central Manager before making any CLI configuration changes on the WAE. Any CLI configuration changes made sooner may be overwritten when the Central Manager updates the WAE. We strongly recommend making all configuration changes by using the Central Manager GUI.

The WAAS CLI is organized into four command modes. Each command mode has its own set of commands to use for the configuration, maintenance, and monitoring of a WAE. The commands that are available to you depend on the mode you are in. When you enter a question mark (?) at the system prompt, you can obtain a list of commands available for each command mode.

The four WAAS command modes are as follows:

- EXEC mode—For setting, viewing, and testing system operations. This mode is divided into two access levels: user and privileged. To use the privileged access level, enter the **enable** command at the user access level prompt, then enter the privileged EXEC password when you see the password prompt.
- Global configuration mode—For setting, viewing, and testing the configuration of WAAS software features for the entire device. To use this mode, enter the **configure** command from the privileged EXEC mode.
- Interface configuration mode—For setting, viewing, and testing the configuration of a specific interface. To use this mode, enter the **interface** command from the global configuration mode.
- Feature-specific configuration mode—Some configuration modes are available from the global configuration mode for managing specific features.

For information about using the CLI to configure a WAAS device, see the *Cisco Wide Area Application Services Command Reference* and the *Cisco Wide Area Application Services Quick Configuration Guide*.

Benefits of Cisco WAAS

This section describes the benefits of Cisco WAAS and includes the following topics:

Preservation of Source TCP/IP Information

Many optimization products create tunnels through routers and other networking devices, which result in a loss of source TCP/IP information in the optimized data. This loss of TCP/IP information often disrupts important network services (such as QoS and NBAR), and can disrupt proper operation of traffic analysis tools such as NetFlow and security products and features such as ACLs and IP-based firewalls.

Unlike other optimization products, Cisco WAAS seamlessly integrates into your network and preserves all TCP/IP header information in the traffic that it optimizes, so that your existing analysis tools and security products are not compromised.

Autodiscovery of WAAS Devices

Cisco WAAS includes an autodiscovery feature that enables WAEs to automatically locate peer WAEs on your network. After autodiscovering a peer device, the WAEs can terminate and separate the LAN-to-WAN TCP connections and add a buffering layer to resolve the differing speeds. Once a WAE establishes a connection to a peer WAE, the two devices can establish an optimized link for TCP traffic, or pass the traffic through as unoptimized.

The autodiscovery of peer WAAS devices is achieved using proprietary TCP options. These TCP options are only recognized and understood by WAAS devices and are ignored by non-WAAS devices.

Centralized Network Monitoring and Management

Cisco WAAS Web-based management tools (WAAS Central Manager and WAE Device Manager GUIs) enable IT administrators to centrally define, monitor, and manage policies for each WAAS device, such as usage quota, backups, disaster recovery, restores, access control, and security policies. IT administrators can also perform the following tasks:

- Remotely provision, configure, and monitor each WAAS device or device group.
- Optimize system performance and utilization with comprehensive statistics, logs, and reporting.
- Perform troubleshooting tasks using tools such as SNMP-based monitoring, traps and alerts, and debug modes.

IT administrators benefit from the following features of Cisco WAAS:

- Native protocol support—Provides complete end-to-end support for the underlying file system protocol (Windows/CIFS) used by the enterprise. Security, concurrency, and coherency are preserved between each client and file server.
- Transparency—Is fully transparent to applications, file systems, and protocols, enabling seamless integration with existing network infrastructures, including mixed environments. Cisco WAAS also has no impact on any security technology currently deployed.
- Branch office data protection—Increases data protection at branch offices. Its file cache appears on the office's LAN in the same way as a local file server. End users can map their personal document folders onto the file cache using Windows or UNIX utilities. A cached copy of user data is stored locally in the branch WAE for fast access. The master copy is stored centrally in the well-protected data center.
- Centralized backup—Consolidates data across the extended enterprise into a data center, which makes it easy to apply centralized storage management procedures to branch office data. Backup and restore operations become simpler, faster, and more reliable than when the data was decentralized.

  In the event of data loss, backup files exist in the data center and can be quickly accessed for recovery purposes. The amount of data loss is reduced because of the increased frequency of backups performed on the centralized storage in the data center. This centralized storage backup makes disaster recovery much more efficient and economical than working with standalone file servers or NAS appliances.
- Simplified storage management—Migrates storage from remote locations to a central data facility, which reduces costs and simplifies storage management for the extended enterprise.
- WAN adaptation—Provides remote users with near-LAN access to files located at the data center. WAAS uses a proprietary protocol that optimizes the way traffic is forwarded between the WAEs.

Optimized Read and Write Caching

The common file services feature in Cisco WAAS maintains files locally, close to the clients. Changes made to files are immediately stored in the local branch WAE, and then streamed to the central file server. Files stored centrally appear as local files to branch users, which improves access performance. CIFS caching includes the following features:

- Local metadata handling and caching—Allows metadata such as file attributes and directory information to be cached and served locally, optimizing user access.
- Partial file caching—Propagates only the segments of the file that have been updated on write requests rather than the entire file.
- Write-back caching—Facilitates efficient write operations by allowing the data center WAE to buffer writes from the branch WAE and to stream updates asynchronously to the file server without risking data integrity.
- Advance file read—Increases performance by allowing a WAE to read the file in advance of user requests when an application is conducting a sequential file read.
- Negative caching—Allows a WAE to store information about missing files to reduce round-trips across the WAN.
- Microsoft Remote Procedure Call (MSRPC) optimization—Uses local request and response caching to reduce the round-trips across the WAN.
- Signaling messages prediction and reduction—Uses algorithms that reduce round-trips over the WAN without loss of semantics.

WCCP Support

The Web Cache Communication Protocol (WCCP) developed by Cisco Systems specifies interactions between one or more routers (or Layer 3 switches) and one or more application appliances, web caches, and caches of other application protocols. The purpose of the interaction is to establish and maintain the transparent redirection of selected types of traffic flowing through a group of routers. The selected traffic is redirected to a group of appliances. Any type of TCP traffic can be redirected.

The WCCP v2 protocol has a built-in set of beneficial features, for example, automatic failover and load balancing. The router monitors the liveness of each WAE attached to it through the WCCP keepalive messages, and if a WAE goes down, the router stops redirecting packets to the WAE. By using WCCP, the branch WAE avoids becoming a single point of failure. The router can also load balance the traffic among a number of branch WAEs.

Cisco WAAS supports transparent interception of TCP sessions through WCCP. Once WCCP is turned on at both the router and the branch WAE, only new sessions are intercepted. Existing sessions are not affected.

PBR Support

Policy-based routing (PBR) allows IT organizations to configure their network devices (a router or a Layer 4 to Layer 6 switch) to selectively route traffic to the next hop based on the classification of the traffic. WAAS administrators can use PBR to transparently integrate a WAE into their existing branch office network and data centers. PBR can be used to establish a route that goes through a WAE for some or all packets based on the defined policies.

For more information about PBR, see Chapter 5, “Configuring Traffic Interception.”

Inline Interception Support

Direct inline traffic interception is supported on WAEs with a Cisco WAE Inline Network Adapter or Interface Module installed. Inline interception of traffic simplifies deployment and avoids the complexity of configuring WCCP or PBR on the routers.

An inline WAE transparently intercepts traffic flowing through it or bridges traffic that does not need to be optimized. It also uses a mechanical fail-safe design that automatically bridges traffic if a power, hardware, or unrecoverable software failure occurs.

Note AppNav Controller Interface Modules do not support automatic bypass mode to continue traffic flow in the event of a failure. For high availability, two or more AppNav Controller Interface Modules should be deployed in an AppNav cluster. For more information on using inline mode with the AppNav solution, see Chapter 4, "Configuring AppNav."

You can configure the inline WAE to accept traffic only from certain VLANs; for all other VLANs, traffic is bridged and not processed.

You can serially cluster inline WAE devices to provide higher availability in the event of a device failure. If the current optimizing device fails, the second inline WAE device in the cluster provides the optimization services. Deploying WAE devices in a serial inline cluster for the purposes of scaling or load balancing is not supported.

For more information about inline mode, see the "Using Inline Mode Interception" section on page 5-42.

Failure Resiliency and Protection

Cisco WAAS provides a high-availability failover (and load-balancing) function that minimizes the probability and duration of CIFS downtime.

If a WAE configured for CIFS fails, all peer WAEs configured to operate with it are redirected to work with an alternate WAE. This operation maintains high availability without service interruption.

This change may not be transparent to users, which means that client connections are closed and require CIFS clients to reestablish their connection. Whether such changes impact currently running applications depends on the behavior of the application being used, and on the behavior of the specific CIFS client. Typically, however, the transition is transparent to the client.

RAID Compatibility

Cisco WAAS provides the following Redundant Array of Independent Disks (RAID) capability for increased storage capacity or increased reliability:

- Logical Disk Handling with RAID-5–Logical disk handling with Redundant Array of Independent Disks-5 (RAID-5) is implemented in WAAS as a hardware feature. RAID-5 devices can create a single logical disk drive that may contain up to six physical hard disk drives, providing increased logical disk capacity.

  Systems with RAID-5 can continue operating if one of the physical drives fails or goes offline.
- Logical Disk Handling with RAID-1—Logical disk handling with RAID-1 is implemented in WAAS as a software feature. RAID-1 uses disk mirroring to write data redundantly to two or more drives, providing increased reliability.

Because the software must perform each disk write operation against two disk drives, the filesystem write performance may be affected.

- Disk Hot-Swap Support—WAAS for RAID-1 allows you to hot-swap the disk hardware. RAID-5 also allows you to hot-swap the disk hardware after the RAID array is shut down. For the disk removal and replacement procedures for RAID systems, see Chapter 16, "Maintaining Your WAAS System."

Streamlined Security

Cisco WAAS supports disk encryption, which addresses the need to securely protect sensitive information that flows through deployed WAAS systems and that is stored in WAAS persistent storage.

Cisco WAAS does not introduce any additional maintenance overhead on already overburdened IT staffs. Cisco WAAS avoids adding its own proprietary user management layer, and instead makes use of the users, user credentials, and access control lists maintained by the file servers. All security-related protocol commands are delegated directly to the source file servers and the source domain controllers. Any user recognized on the domain and source file server are automatically recognized by Cisco WAAS with the same security level, and all without additional configuration or management.

Cisco WAAS delegates access control and authentication decisions to the origin file server.

SNMP Support

Cisco WAAS supports Simple Network Management Protocol (SNMP) including SNMPv1, SNMPv2, and SNMPv3. Cisco WAAS supports many of the most commonly used SNMP managers, such as HP OpenView and IBM Tivoli NetView.

Most Cisco WAAS traps are also recorded in the logs displayed in the WAAS Central Manager GUI, although some (such as exceeding the maximum number of sessions) are reported only to the SNMP manager.

Cisco WAAS supports parameters based on SNMPv2, enabling it to integrate into a common SNMP management system. These parameters enable system administrators to monitor the current state of the WAAS network and its level of performance.

Exported parameters are divided into the following categories:

- General parameters—Includes the version and build numbers and license information.
- Management parameters—Includes the location of the Central Manager.
- Data center WAE parameters—Includes the general parameters, network connectivity parameters, and file servers being exported.
- Branch WAE parameters—Includes the general parameters, network connectivity parameters, CIFS statistics, and cache statistics.

For more information about SNMP and supported MIBs, see Chapter 18, "Configuring SNMP Monitoring."

CHAPTER 2

Planning Your WAAS Network

This chapter describes general guidelines, restrictions, and limitations that you should be aware of before you set up your Wide Area Application Services (WAAS) network.

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE and WAVE appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS, and vWAAS instances.

This chapter contains the following sections:

Checklist for Planning Your WAAS Network

Cisco Wide Area Application Engines (WAEs) that are running the WAAS software can be used by enterprises or service providers to optimize the application traffic flows between their branch offices and data centers. You deploy WAE nodes at the WAN endpoints near the networked application clients and their servers, where they intercept WAN-bounded application traffic and optimize it. You must insert WAE nodes into the network flow at defined processing points.

WAAS software supports the following three typical network topologies:

- Hub and spoke deployments—In a hub and spoke deployment servers are centralized and branch offices host clients and a few local services only (for example, WAAS printing services).

- Mesh deployments—In a mesh deployment, any location may host both clients and servers and the clients may access any number of local or remote servers.
- Hierarchical deployments—In a hierarchical deployment, the servers are located in multiple regional, national data centers and are accessed by the different clients. The connections between the data centers are of higher bandwidth than the connections to the branch offices.

The deployments are characterized according to the WAAS element connections, which follow the client-server access pattern and may differ from the physical network links. For more information, see Chapter 1, "Introduction to Cisco WAAS."

Planning Checklist

When you are planning your WAAS network, use the following checklist as a guideline. As the following checklist indicates, you can break the planning phase into the following three main categories of planning activities:

- Sizing phase
- Planning for management
- Planning for application optimization

Note Although there are some interdependencies, you do not need to complete all of the steps in a particular planning phase before you start the next step.

To plan your network, follow these guidelines:

1. Complete the sizing phase that includes the following tasks:

- Determine which locations in your existing network require WAAS optimization (for example, which branch offices and data centers).
- Determine if you are going to use a traditional WAAS deployment model or the AppNav deployment model. For more information on AppNav, see Chapter 4, "Configuring AppNav."
- Determine the number and models of the WAAS devices that are required for each location. Some key factors in this selection process is the WAN bandwidth, the number of users, and the expected use. Various hardware configurations are possible (for example, different hard disk models and RAM size). Consider running a cluster of WAEs where additional scalability and or failover is required. For more information, see the "Calculating the Number of WAAS Devices Needed" section on page 2-18.
- Verify that you have purchased sufficient licenses to cover your needs.

2. Plan for management as follows:

- Complete site and network planning (for example, obtain the IP and routing information including IP addresses and subnets, routers and default gateway IP addresses, and the hostnames for the devices). See the "Checklist of WAAS Network System Parameters" table in the *Cisco Wide Area Application Services Quick Configuration Guide*.
- Determine the login authentication and login authorization methods (for example, external RADIUS, TACACS+, Windows domain servers) and accounting policies that you want your WAAS Central Managers and WAEs to use. For more information, see Chapter 7, "Configuring Administrative Login Authentication, Authorization, and Accounting."

- For security purposes, plan to change the predefined password for the predefined superuser account immediately after you have completed the initial configuration of a WAE. For more information, see "WAAS Login Authentication and Authorization" section on page 2-26.
- Determine if you need to create any additional administrative accounts for a WAAS device. For more information, see Chapter 8, "Creating and Managing Administrator User Accounts and Groups."
- Determine if you should group your WAEs into logical groups. For more information, see the "Logically Grouping Your WAEs" section on page 2-27.
- Determine which management access method to use. By default, Telnet is used but SSH may be the preferred method in certain deployments. For more information, see the "Configuring Login Access Control Settings for WAAS Devices" section on page 7-7.

3. Plan for application optimization as follows:

- Determine and resolve router interoperability issues (for example, the supported hardware and software versions, router performance with interception enabled). For more information, see the "Site and Network Planning" section on page 2-4.
- Determine the appropriate interception location when the data center or branch office is complex (for example, if your existing network uses a hierarchical topology).
- Determine which WAAS services to deploy. For more information about the different WAAS services, see Chapter 1, "Introduction to Cisco WAAS."
- Determine which WAAS software licenses to install. Software licenses enable specific WAAS services. For more information about installing software licenses, see the "Managing Software Licenses" section on page 10-3.
- Determine which traffic interception methods to use in your WAAS network (for example, inline mode, WCCP Version 2, or policy-based routing (PBR)). For more information, see the "Supported Methods of Traffic Redirection" section on page 2-19.

  Note WCCP works only with IPv4 networks.

- If you plan to use the WCCP TCP promiscuous mode service as a traffic interception method, determine whether you should use IP access control lists (ACLs) on your routers.

  Note IP ACLs that are defined on a router take precedence over the ACLs that are defined on the WAE. For more information, see the "Access Lists on Routers and WAEs" section on page 2-25.

- Determine whether you need to define IP ACLs or interception ACLs on the WAEs. For more information, see the "Access Lists on Routers and WAEs" section on page 2-25.

  Note ACLs that are defined on a WAE take precedence over the WAAS application definition policies that are defined on the WAE.

- If PBR is to be used, determine which PBR method to use to verify PBR next-hop availability for your WAEs. For more information, see the "Methods of Verifying PBR Next-Hop Availability" section on page 5-39.
- Determine the major applications for your WAAS network. Verify whether the predefined application definition policies cover these applications and whether you should add policies if your applications are not covered by these predefined policies. For a list of the predefined application

definition policies, see Appendix A, "Predefined Optimization Policy."

- Consider day zero migration of file systems if file servers are to be centralized in the process. For more information, see the "Data Migration Process" section on page 2-28.

After you complete the planning tasks, you are ready to perform a basic configuration of a WAAS network as described in the *Cisco Wide Area Application Services Quick Configuration Guide*.

Site and Network Planning

Before you install and deploy WAAS devices in your network, you need to collect information about your network to accommodate the integration of the WAAS devices.

In a typical distributed organizational layout, there are two types of networks where WAAS devices are installed:

- The data center (central office), where one or more colocated data center WAEs provide access to the resident file and application servers. In data centers, you can deploy a WAE as a single device or a pair of WAEs as a high-availability or load-sharing pairs. High-availability pairs are supported if either WCCP Version 2 or PBR is being used for traffic redirection in the data center; load-sharing pairs are only supported if WCCP Version 2 is being used for traffic redirection in the data center.
- The branch offices, where branch WAEs enable users to access the file and application servers over the WAN. In branch offices, you can deploy a WAE as a single device or a pair of WAEs as a high-availability or load-sharing pairs. High-availability pairs are supported if either WCCP Version 2 or PBR is being used for traffic redirection in the branch office; load-sharing pairs are only supported if WCCP Version 2 is being used for traffic redirection in the branch office.

In collaborative networks, colocated data center WAEs and branch WAEs are deployed throughout the network. These colocated WAEs are configured to share data in opposite directions (two cross-linked servers).

The WAE attaches to the LAN as an appliance. A WAE relies on packet interception and redirection to enable application acceleration and WAN optimization. Consequently, traffic interception and redirection to a WAE must occur at each site where a WAE is deployed. Traffic interception and redirection occurs in both directions of the packet flow. Because Layer 3 and Layer 4 headers are preserved, you may need to ensure that you always connect a WAE to a tertiary interface (or a subinterface) on the router to avoid routing loops between the WAE and WCCP or PBR-enabled router that is redirecting traffic to it. For more information on this topic, see the "Using Tertiary Interfaces or Subinterfaces to Connect WAEs to Routers" procedure on page 2-24.

Note We strongly recommend that you do not use half-duplex connections on the WAE or on routers, switches, or other devices. Half duplex impedes performance and should not be used. Check each Cisco WAE interface and the port configuration on the adjacent device (router, switch, firewall, or WAE) to verify that full duplex is configured.

Note The data center WAE and branch WAE communicate with each other only if the firewall is open.

Note WAAS versions 5.x and lower do not support timestamp (TSVAL) TCP option.

This section contains the following topics:

Windows Network Integration

To successfully integrate WAAS devices into the Windows environment, you might need to make certain preparations on both the data center WAE and branch WAE sides of the network. This section contains the following topics:

Data Center WAE Integration

Before the initial configuration of the data center WAE, you need to know the following parameters:

- WINS server (if applicable).
- DNS server and DNS domain (if applicable).
- A browsing user with file-server directory traversal (read-only) privileges. This user, who is usually set up as a domain or service user, is required for running preposition policies.

To successfully integrate Cisco WAAS into the Windows environment on the data center WAE side of a network where DHCP is not being used, you must manually add the name and IP address of the data center WAE to the DNS server. You should take this action before installing and deploying the WAAS devices.

Note User permissions are determined by the existing security infrastructure.

Branch WAE Integration

Before the initial configuration of the branch WAE, you need to know the following parameters:

- DNS server and DNS domain
- Windows Domain Name
- WINS server (if applicable)

To successfully integrate Cisco WAAS into the Windows environment on the branch WAE side of the network, you should take the following preliminary actions before installing and deploying the WAAS devices in your network:

- To enable all branch WAEs in the specified domain to appear in the Network Neighborhood of users within the same domain, ensure that a Domain Master Browser or local Master Browser is active.
- If DHCP is not used, you must manually add the name and IP address of the branch WAE to the DNS server.

UNIX Network Integration

Before the initial configuration of a WAAS device, you need to know the following parameters:

- DNS server and DNS domain.
- NIS server parameters (if applicable).
- On the data center WAE side, a browsing UID or GID with file-server directory traversal (read-only) privileges. This UID or GID, which is usually set up as a domain or service user, is required for browsing when defining coherency policies.

To successfully integrate Cisco WAAS into the UNIX environment, you need to perform these actions on both the data center WAE and branch WAE sides of the network:

- You must manually add the name and IP address of both the data center WAE and the branch WAE to the DNS server.
- When separate domains are used, UNIX users may be defined at the remote (branch) offices or on the central servers. This situation may result in the same user name being defined in different domains. A user may be defined differently in the branch and center or may be defined only on one end and not on the other. You can ensure consistency in such cases by using NIS or by mapping between the different domains, either manually or automatically. That is, users can be mapped from the remote server to the central servers by translating their identities from the central office to the remote offices.

Note To map users using automatic management, you must first configure the NIS server in both the data center WAE (primary) and branch WAE (secondary).

CIFS-Related Ports in a WAAS Environment

This section describes the CIFS-related ports used between your clients, WAEs that are functioning as file engines, and CIFS file servers. Most CIFS communication occurs between the branches and the central office. This communication is encrypted and delivered through the organization's VPN. No ports on the firewall need to be opened because all communication is tunneled internally.

You only need to change the firewall setup if administrative or other maintenance work needs to be done from a location outside the organization.

Ports 139 and 445

If you have only deployed CIFS services in your WAAS network, your WAAS network uses ports 139 and 445 to connect clients to a branch WAE and to connect a data center WAE to the associated file servers. The port used depends on the configuration of your WAAS network.

If WCCP is enabled or inline mode is used, the branch WAE accepts client connections on ports 139 or 445. If neither WCCP nor inline mode are enabled, the branch WAE accepts connections only over port 139.

Your WAAS network always tries to use the same port to communicate end-to-end. Consequently, if a client uses port 445 to connect to a branch WAE, the associated data center WAE will try to use the same port to connect to the file server. If port 445 is unavailable, the data center WAE will try to use port 139.

Some organizations close port 139 on their networks to minimize security risks associated with this port. If your organization has closed port 139 for security reasons, you can configure your WAAS network to bypass port 139. If this is the case in your organization, you need to perform the following task to bypass port 139 and use port 445 in its place if you have only deployed the CIFS services in your WAAS network:

- Enable WCCP Version 2 on your routers and branch WAE, as described in the *Cisco Wide Area Application Services Quick Configuration Guide*. Alternatively, you can use inline mode on a branch WAE with a Cisco WAE Inline Network Adapter or Cisco Interface Module installed.

Ports 88 and 464

If you are using Windows Domain authentication with Kerberos enabled, the WAE uses ports 88 and 464 to authenticate clients with the domain controller.

Firewalls and Directed Mode

By default, WAAS transparently sets up new TCP connections to peer WAEs, which can cause firewall traversal issues when a WAAS device tries to optimize the traffic. If a WAE device is behind a firewall that prevents traffic optimization, you can use the directed mode of communicating to a peer WAE. In directed mode, all TCP traffic that is sent to a peer WAE is encapsulated in UDP, which allows a firewall to either bypass the traffic or inspect the traffic (by adding a UDP inspection rule).

Any firewall between two WAE peers must be configured to pass UDP traffic on port 4050, or whatever custom port is configured for directed mode if a port other than the default is used.

If a WAE using directed mode is behind a NAT device, you must configure the NATed IP address on the WAE.

For more information about configuring directed mode, see the "Configuring Directed Mode" section on page 6-27.

Firewalls and Standby Central Managers

Primary and standby Central Managers communicate on port 8443. If your network includes a firewall between primary and standby Central Managers, you must configure the firewall to allow traffic on port 8443 so that the Central Managers can communicate and stay synchronized.

Performance Tuning for High WAN Bandwidth Branch Offices

WAAS combines Layer-4 TCP optimizations with Layer-7 application accelerators for various protocols including CIFS. For some branch offices with high WAN bandwidth (for example, above 50 Mbps), if the native latency is low (for example, below 20 ms RTT), depending on the number of user sessions and data patterns, applying Layer-4 optimizations alone may provide optimal levels of performance. In such cases, we recommend measuring end-user response times under production load to determine the appropriate operational state for the application accelerators and sizing.

About Autoregistration and WAEs

Autoregistration automatically configures network settings and registers WAEs with the WAAS Central Manager device. On startup, devices running WAAS software (with the exception of the WAAS Central Manager device itself) automatically discover the WAAS Central Manager device and register with it. You do not need to manually configure the device. This feature is useful for large scale automated deployments of devices. Once a WAE is registered, you configure the device remotely using the WAAS Central Manager GUI.

In the example configuration provided in the *Cisco Wide Area Application Services Quick Configuration Guide,* the autoregistration feature is disabled on the WAEs when the setup utility is used to perform the initial configuration of the device.

Autoregistration uses a form of Dynamic Host Configuration Protocol (DHCP). For autoregistration to function, you must have a DHCP server that is configured with the hostname of the WAAS Central Manager and that is capable of handling vendor class option 43.

Note The form of DHCP used for autoregistration is *not* the same as the interface-level DHCP that is configurable through the **ip address dhcp** interface configuration command. (For a description of the **ip address dhcp** interface configuration command, see the *Cisco Wide Area Application Services Command Reference.*)

The vendor class option (option 43) information needs to be sent to the WAAS device in the format for encapsulated vendor-specific options as provided in RFC 2132. The relevant section of RFC 2132, Section 8.4, is reproduced here as follows:

> The encapsulated vendor-specific options field should be encoded as a sequence of code/length/value fields of syntax identical to that of the DHCP options field with the following exceptions:
>
> **a.** There should not be a “magic cookie” field in the encapsulated vendor-specific extensions field.
>
> **b.** Codes other than 0 or 255 may be redefined by the vendor within the encapsulated vendor-specific extensions field but should conform to the tag-length-value syntax defined in section 2.
>
> **c.** Code 255 (END), if present, signifies the end of the encapsulated vendor extensions, not the end of the vendor extensions field. If no code 255 is present, then the end of the enclosing vendor-specific information field is taken as the end of the encapsulated vendor-specific extensions field.

In accordance with the RFC standard, the DHCP server needs to send the WAAS Central Manager’s hostname information in code/length/value format (code and length are single octets). The code for the WAAS Central Manager’s hostname is 0x01. DHCP server management and configuration are not within the scope of the autoregistration feature.

Note The WAE sends “CISCOCDN” as the vendor class identifier in option 60 to facilitate your grouping of WAEs into device groups.

Autoregistration DHCP also requires that the following options be present in the DHCP server’s offer to be considered valid:

- Subnet-mask (option 1)
- Routers (option 3)

- Domain-name (option 15)
- Domain-name-servers (option 6)
- Host-name (option 12)

In contrast, interface-level DHCP requires only subnet-mask (option 1) and routers (option 3) for an offer to be considered valid; domain-name (option 15), domain-name-servers (option 6), and host-name (option 12) are optional. All of the above options, with the exception of domain-name-servers (option 6), replace the existing configuration on the system. The domain-name-servers option is added to the existing list of name servers with the restriction of a maximum of eight name servers.

Autoregistration is enabled by default on the first interface of the device. On an NME-WAE module, autoregistration is enabled on the configured interface. On an SM-SRE module, autoregistration is disabled by default.

Note You must disable autoregistration when both device interfaces are configured as port-channel interfaces.

If you do not have a DHCP server, the device is unable to complete autoregistration and eventually times out. You can disable autoregistration at any time after the device has booted and proceed with manual setup and registration.

To disable autoregistration, or to configure autoregistration on a different interface, use the **no auto-register enable** command in global configuration mode.

Note Autoregistration is automatically disabled if a static IP address is configured or if interface-level DHCP is configured on the same interface as autoregistration. (See the "Selecting Static IP Addresses or Using Interface-Level DHCP" section on page 2-9.)

The following example disables autoregistration on the interface GigabitEthernet 1/0:

```
WAE(config)# no auto-register enable GigabitEthernet 1/0
```

Autoregistration status can be obtained by using the following **show** EXEC command:

```
WAE# show auto-register
```

Selecting Static IP Addresses or Using Interface-Level DHCP

During the initial configuration, you have the option of configuring a static IP address for the device or choosing DHCP.

DHCP is a communications protocol that allows network administrators to manage their networks centrally and automate the assignment of IP addresses in an organization's network. When an organization sets up its computer users with a connection to the network, an IP address must be assigned to each device. Without DHCP, the IP address must be entered manually for each computer, and if computers move to another location in another part of the network, the IP address must be changed accordingly. DHCP automatically sends a new IP address when a computer is connected to a different site in the network.

If you have a DHCP server configured, autoregistration will automatically configure the network settings and register WAEs with the WAAS Central Manager device upon bootup.

If you do not have a DHCP server configured, or you have a DCHP server but do not want to use the autoregistration feature, then manually configure the following network settings with the interactive setup utility or CLI, then register the WAEs with the WAAS Central Manager device. Configure these settings:

- Ethernet interface
- IP domain name
- Hostname
- IP name server
- Default gateway
- Primary interface

When a WAAS device boots, you are prompted to run the first-time setup utility (enter basic configuration), which you use to set up the basic device network settings for the WAE.

Identifying and Resolving Interoperability Issues

This section describes how to identify and resolve interoperability issues. It contains the following topics:

Interoperability and Support

This section contains the following topics:

For a list of the hardware, CIFS clients, and web browsers supported by the WAAS software, see the *Release Note for Cisco Wide Area Application Services.*

Unicode Support for the WAAS GUI Interfaces

The WAAS software supports Unicode in the WAAS Central Manager and the WAE Device Manager GUI interfaces.

In the WAAS Central Manager, you can create preposition policies that include Unicode characters. For example, you can define a preposition policy for a directory that contains Unicode characters in its name.

Specifically, the following fields in the WAAS Central Manager GUI support Unicode:

- The root directory and file pattern fields in the preposition policies

In the WAE Device Manager GUI, you can include Unicode characters in the name of the backup configuration file. In addition, the logs included in the WAE Device Manager GUI can display Unicode characters.

Unicode Support Limitations

The following are Unicode support limitations:

- Usernames cannot contain Unicode characters.
- When defining policies for coherency, and so on, you cannot use Unicode characters in the Description field.
- File server names cannot contain Unicode characters.

WAAS and Cisco IOS Interoperability

This section describes the interoperability of the WAAS software with the Cisco IOS features for a basic WAAS deployment that uses WCCP-based interception and transparent transport and contains the following topics:

Note The WAAS software does not support Cisco IOS IPv6 and Mobile IP.

We recommend that you use Cisco IOS Software Release 12.2 or later.

WAAS Support of the Cisco IOS QoS Classification Feature

You classify packets by using a policy filter (for example, using QPM) that is defined on the packets. You may use the following policy filter properties:

- Source IP address or hostname—Supported under WAAS because the source IP address is preserved by the WAAS device.
- Source TCP/UDP port (or port range)—Supported under WAAS because the source port is preserved by the WAAS device.
- Destination IP address or hostname—Supported under WAAS because the destination IP is preserved by WAAS. WAAS relies on interception at the data center for redirecting traffic to the peer WAAS device.
- Destination TCP/UDP port (or port range)—Supported under WAAS because the destination IP is preserved by WAAS. WAAS relies on interception at the data center for redirecting traffic to the peer WAAS device.

- DSCP/IP precedence (TOS)—Supported under WAAS because WAAS copies the settings of incoming packets on to the outgoing packets from WAAS back to the router. If the packets are not colored at connection establishment time (for TCP packets), there might be a delay in propagating the settings because WAAS does not poll these settings periodically. The packets are eventually colored properly. When packets are not colored they are left uncolored by the WAAS software.

WAAS software does not support IPv6 QoS, MPLS QoS, ATM QoS, Frame Relay QoS, and Layer 2 (VLAN) QoS.

WAAS Support of the Cisco IOS NBAR Feature

Unlike a traditional type of classification that is specified through a policy filter that is listed in the "WAAS Support of the Cisco IOS QoS Classification Feature" section on page 2-11, Network-Based Application Recognition (NBAR) classification needs to consider payload. The classification keeps track of any interceptor that modifies the payload because this modification might cause NBAR to not be able to classify the packets. However, the WAAS software does support NBAR.

The following is an example flow of how the WAAS software supports NBAR:

1. A packet P1, which is part of a TCP stream S1, enters the router and is classified by NBAR on the LAN interface of the router as belonging to class C1. If the classification of P1 does not involve payload inspection (for example, only TCP/IP headers), no action needs to be taken because the WAAS software preserves this information.
2. If P1 classification requires payload inspection, P1 needs to be marked using the TOS/DSCP bits in the packet (as opposed to using other internal marking mechanisms).
3. P1 is then intercepted through WCCP Version 2 (still on the LAN interface, WCCP is processed after NBAR) and is redirected to a WAE.
4. WAAS applies any optimizations on the payload and copies the DCSP bits settings from the incoming TCP stream, S1 onto the outgoing stream, S2 (which is established between the local WAAS appliance and the remote WAAS appliance over the WAN). Because NBAR usually needs to see some payload before doing the classification, it is unlikely that WAAS will have the proper bit settings at connection-establishment time. Consequently, the WAAS software uses polling to inspect the DSCP bits on the incoming TCP stream, then copies it over to the stream from the WAAS device back to the router.
5. When S2 reenters the router, NBAR will not classify S2 as belonging to C1 because the payload has been changed or compressed. However, the DSCP settings have already marked these packets as belonging to C1. Consequently, these packets will be treated properly as if they were classified through NBAR.

   As long as the flow is not identified, NBAR will continue to search for classification in the packets. Because compressed packets will not be classified, this situation can unnecessarily burden the CPU (doing packet inspection). Because of the potential degradation in performance and the slight possibility of correctness issues, we strongly recommend that you use a subinterface or a separate physical interface to connect the WAE to the router (as described in the "Using Tertiary Interfaces or Subinterfaces to Connect WAEs to Routers" section on page 2-24). When you use a tertiary interface or subinterface to connect the WAE to the router, both the performance and correctness issues are addressed because each packet is processed only once.
6. For dynamic classifications, NBAR maintains a per-flow state. Once certain flows are classified, NBAR does not continue to perform deep packet inspection anymore. However, for other flows (for example, Citrix), NBAR does look at packets continuously because the classification may change dynamically in a flow. Therefore, in order to support all NBAR classifications, it is not sufficient to only poll the DSCP settings of packets incoming to WAAS once per flow; you need to poll

periodically to identify flow changes. However, the WAAS system expects packets to appear in the sequence of packets belonging to class C1, followed by a sequence of C2, and so forth, so that a polling method is sufficient to track such dynamic changes.

Note This dynamic classification support requires support for marking DSCP/ToS settings, as specified in the "WAAS Support of the Cisco IOS QoS Classification Feature" section on page 2-11, as well as the tracking of dynamic changes through polling.

Several router configurations need to be followed in order to ensure NBAR-WAAS compliance, and you must ensure that the following router configurations are adhered to:

- Ensure that classification is followed by proper DSCP marking.
- Ensure that the router in general (IP access lists that are configured on the router) does not scrub DSCP/TOS settings that are already marked on the packets on entry, and that NBAR does not unmark marked packets.

WAAS Support of the Cisco IOS Marking

The Cisco IOS marking feature is supported by the WAAS software.

WAAS Support of the Cisco IOS Queuing

The Cisco IOS queuing feature for congestion management is supported by the WAAS software.

WAAS Support of the Cisco IOS Congestion Avoidance

The Cisco IOS congestion avoidance feature is supported by the WAAS software.

WAAS Support of the Cisco IOS Traffic Policing and Rate Limiting

The Cisco IOS traffic policing and rate-limiting feature is only partially supported by the WAAS software. This Cisco IOS feature will work properly when enabled on an outbound interface. However, when this feature is enabled on an inbound interface, it will see both compressed and uncompressed traffic, and will result in inaccurate rate limiting.

WAAS Support of the Cisco IOS Signaling

The Cisco IOS signaling (RSVP) feature is typically implemented in MPLS networks. Because the WAAS software does not interact with MPLS RSVP messages, the RSVP feature is supported.

WAAS Support of the Cisco IOS Link-Efficiency Operations

The Cisco IOS link-efficiency operations are supported by the WAAS software.

WAAS Support of the Cisco IOS Provisioning, Monitoring, and Management

The Cisco IOS AutoQoS feature is supported by the WAAS software but requires additional configuration. This feature is closely connected with NBAR support because the AutoQoS feature uses NBAR to discover the various flows on the network. However, because the Cisco IOS AutoQoS feature is strictly on an outbound feature (for example, it cannot be enabled on the inbound side of an interface), this situation could create a potential problem because enabling NBAR on the outbound interface is not supported.

To avoid this potential problem, enable the trust option of the AutoQoS feature on the following interfaces so that classification and queuing are performed based on the marked value (NBAR is not enabled on the outbound interface using this solution):

- On the LAN interface on which the input policy is created and on which the marking of the packets should be performed according to the AutoQoS marking (for example, interactive video mark to af41).
- On the WAN outbound interface.

WAAS and Management Instrumentation

For management instrumentation use with the WAAS software, note the following:

- When deployed in native (transparent) mode, WAAS maintains packet header information vital to technologies such as NetFlow. NetFlow can be configured on adjacent devices and exports flow record information in accordance with where NetFlow is configured in relation to the WAAS device. For NetFlow configurations on the LAN side of a WAAS device, NetFlow exports records containing information about original flows. For NetFlow configurations on the WAN side of a WAAS device, NetFlow exports records containing information about optimized and pass-through flows.
- You may see statistics on optimized and unoptimized traffic.
- IP Service Level Agreements (SLAs) are supported.
- Full support of policies based on Layer 3 and Layer 4 is provided. Policies based on Layer 7 are partially supported because the first few messages are unoptimized.
- Intrusion Detection System (IDS) is partially supported. The first few messages are unoptimized to allow IDS to detect the intrusive strings.
- Cisco IOS security is partially supported with the exception of features that rely on Layer 5 and above visibility.
- IPsec and SSL VPN is supported.
- Access control lists (ACLs) are supported. IP ACLs on the router take precedence over ACLs that are defined on the WAE. For more information, see the "Access Lists on Routers and WAEs" section on page 2-25.
- VPN is supported if the VPN is deployed after WCCP interception occurs.

Note A WAAS device does not encrypt WAN traffic. If you require additional security measures, you should use a VPN. However, the VPN appliances must encrypt and decrypt traffic after and before the WAAS devices so that the WAAS device only sees unencrypted traffic. The WAAS device is unable to compress encrypted traffic and provides only limited TCP optimization to it.

- Network Address Translation (NAT) is supported. However, payload-based NAT is not supported.

WAAS and MPLS

MPLS is partially supported by the WAAS software. WCCP does not know how to operate with packets that are tagged with MPLS labels. Consequently, inside the cloud, WCCP redirection will not function (for example, WCCP redirection will not work for intermediate WAEs). However, as long as the redirection occurs on interfaces that are outside the MPLS cloud, WAAS is supported.

WAAS Compatibility with other Cisco Appliances and Software

If a firewall is placed between the clients and the WAE on one side, and the router on the other side of the firewall, default WCCP redirection does not work. However, if there is a router inside the firewall and another router outside the firewall, default WCCP-based redirection does work and WAAS is supported. You can also enable directed mode to avoid firewall traversal issues. For more information, see the "Configuring Directed Mode" section on page 6-27.

Support for concatenating ACNS and WAAS devices in your network is supported. ACNS devices optimize web protocols and can be used to serve content locally. WAAS devices optimize requests from a Content Engine, which is an ACNS device that needs service from an upstream server or an upstream Content Engine. The ability to concatenate ACNS and WAAS devices in a network has the following benefits:

- If you have already deployed ACNS in your network, you can also deploy WAAS.
- If you have not already deployed ACNS in your network, but need certain ACNS features, you can purchase ACNS and deploy it with WAAS.

WAAS Devices and Device Mode

You must deploy the WAAS Central Manager on a dedicated appliance. Although the WAAS Central Manager device runs the WAAS software, its only purpose is to provide management functions. WAAS Central Manager communicates with the WAEs, which are registered with it, in the network. Through the WAAS Central Manager GUI, you can centrally manage the configuration of the WAEs individually or in groups. WAAS Central Manager also gathers management statistics and logs for its registered WAEs.

A WAE also runs the WAAS software, but its role is to act as an accelerator in the WAAS network.

In a WAAS network, you must deploy a WAAS device in one of the following device modes:

- WAAS Central Manager mode—Mode that the WAAS Central Manager uses.
- WAAS application accelerator mode—Mode that a WAAS Accelerator (data center WAEs and branch WAEs that run the WAAS software) uses to optimize and accelerate traffic.
- WAAS AppNav Controller mode—Mode for a WAAS device that is operating as an AppNav Controller (ANC) that is intercepting and distributing traffic to other WAAS devices operating in application accelerator mode.

The default device mode for a WAAS device is WAAS accelerator mode. The **device mode** global configuration command allows you to change the device mode of a WAAS device.

For example, after you use the WAAS CLI to specify the basic network parameters for the designated WAAS Central Manager (the WAAS device named waas-cm) and assign it a primary interface, you can use the **device mode** configuration command to specify its device mode as central-manager.

```
waas-cm# configure
waas-cm(config)# primary-interface gigabitEthernet 1/0
waas-cm(config)# device mode central-manager
waas-cm(config)# exit
waas-cm# copy run start
waas-cm# reload
Proceed with reload?[confirm]yes
Shutting down all services, will timeout in 15 minutes.
reload in progress ..
```

For more information about how to initially configure a WAAS device, see the *Cisco Wide Area Application Services Quick Configuration Guide*.

Note You cannot configure a WAE network module in the NME-WAE or SM-SRE family of devices to operate in WAAS Central Manager mode.

You can configure a WAE with a Cisco WAE Inline Network Adapter to operate in WAAS Central Manager mode, but the inline interception functionality is not available.

Changing Device Mode

If you want to change the device mode of a device that is already registered with a Central Manager, you must first deregister the device from the Central Manager, change the device mode, reload the device, and then reenable CMS services.

The following steps show how to change the device mode from application-accelerator to appnav-controller:

Step 1 Deregister the device from the Central Manager.

```
wae# cms deregister
Deregistering WAE device from Central Manager will result in loss of data on encrypted
file systems, imported certificate/private keys for SSL service and cifs/wafs preposition
credentials. If secure store is initialized and open, clear secure store and wait for one
datafeed poll rate to retain cifs/wafs preposition credentails.
Do you really want to continue (yes|no) [no]?yes
Disabling management service.
management services stopped
Sending de-registration request to CM
SSMGR RETURNING: 7 (Success)
Removing cms database tables.
Re-initializing SSL managed store and restarting SSL accelerator.Deregistration complete.
Save current cli configuration using 'copy running-config startup-config' command because
CMS service has been disabled.
```

Step 2 Change the device mode to appnav-controller.

```
wae# configure
wae(config)# device mode appnav-controller
The new configuration will take effect after reload.
```

Step 3 Save the configuration and reload.

```
wae(config)# exit
wae# copy run start
wae# reload
Proceed with reload?[confirm]yes
Proceed with clean WCCP shutdown?[confirm]yes

WCCP clean shutdown initiated
Waiting for shutdown ok (1 seconds) . Press ^C to skip waiting
WCCP clean shutdown wait time expired
Shutting down all services, will timeout in 15 minutes.
reload in progress ..
```

Step 4 Log into the WAE after it finished rebooting.

```
AppNav Controller


wae login: admin
Password:
System Initialization Finished.
wae#
```

Step 5 Reenable CMS services.

```
wae# config
wae(config)# cms enable
Registering WAAS AppNav Controller...
Sending  device registration request to Central Manager with address 10.43.65.50
Please wait, initializing CMS tables
Successfully initialized CMS tables
Registration complete.
Please preserve running configuration using 'copy running-config startup-config'.
Otherwise management service will not be started on reload and node will be shown
'offline' in WAAS Central Manager UI.
management services enabled
```

Step 6 Save the configuration.

```
wae(config)# exit
wae# copy run start
```

Note While using the AppNav IOM on the WAVE devices (7571 and 8541 only); and when the device mode is AppNav Controller, the total TCP connection capacity for optimized traffic is reduced. This affects only the local device hosting the AppNav Controller and not the connection capacity or the throughput of the AppNav Controller itself. There is no impact on any other device that is a part of the AppNav cluster.
The table shows the reduced number of connections on the AppNav Controller.

Platform	Application Accelerator	Appnav-mode
7571	60,000	50,000
8541	1,50,000	1,40,000

Calculating the Number of WAAS Devices Needed

When the threshold value of an operational system aspect is exceeded, Cisco WAAS may not meet its expected service level. This situation might result in degraded performance.

The source of the limitation might originate from a specific Cisco WAAS device (WAAS Central Manager, branch WAE, or data center WAE), the entire Cisco WAAS system, a hardware constraint, or the network connecting the distributed software entities. In some cases, the limitation might be resolved by adding more resources or by upgrading the hardware or software.

When planning your network, consider the operational capacity, such as the number of users it should support, how many files it should support, and how much data it should cache.

When planning your WAAS network, refer to the following additional guidelines:

- Number of WAAS Central Managers— All networks must have at least one WAAS Central Manager. For larger networks, you should consider deploying two WAAS Central Managers for active and standby back-up, high availability, and failover. A WAAS Central Manager is deployed on a dedicated appliance.
- Number of WAEs—A minimum of two WAEs are required for traffic optimization; one WAE is required on either side of a network link (for example, one in the branch office and one in the data center). A single site can have more than one WAE for redundancy purposes.
- Number of branch WAEs—At least one branch WAE is required in each remote office. Larger offices usually have multiple departments whose users work with different servers in the central office. In this situation, you can manage your system easier by following the organizational structure with a branch WAE for each department. In certain situations, multiple branch WAEs can be clustered and configured using WCCP to provide failover capabilities. WCCP is the recommended method for larger user populations.
- Number of data center WAEs—Each organization must have at least one data center WAE.
- Number of ANCs—If you are using the AppNav deployment model, at least one ANC is required.

When determining the number of the component types required by your organization, consider the following factors:

- Number of users connecting to the system—This number depends on the static and dynamic capacities defined for the system:
  - Static capacities—Defines the number of user sessions that can connect to the system before it reaches its capacity.
  - Dynamic capacities—Defines the amount of traffic handled by the servers, which means the amount of work being performed on the network. For example, consider whether the users currently connected to the system place a heavy or light load on it.

Note You should calculate dynamic limits based on the specific load assumptions that are particular to each customer.

- Total number of users in all branches that connect to the file servers through the data center WAE—When the number of users is more than one data center WAE can support, you must add one or more additional data center WAEs to the network.

Supported Methods of Traffic Redirection

In a WAAS network, traffic between the clients in the branch offices and the servers in the data center can be redirected to WAEs for optimization, redundancy elimination, and compression. Traffic is intercepted and redirected to WAEs based on policies that have been configured on the routers. The network elements that transparently redirect requests to a local WAE can be a router using WCCP version 2 or PBR to transparently redirect traffic to the local WAE or a Layer 4 to Layer 7 switch (for example, the Catalyst 6500 series Content Switching Module [CSM] or Application Control Engine [ACE]).

Alternately, a WAE that has the Cisco WAE Inline Network Adapter or Cisco Interface Module installed can operate in inline mode and receive and optimize traffic directly before it passes through the router.

In an AppNav deployment, an AppNav Controller in the data center receives intercepted traffic through WCCP, PBR, or inline mode and distributes it to WAAS nodes that optimize the traffic. For more information on an AppNav deployment, see Chapter 4, "Configuring AppNav."

This section contains the following topics:

For detailed information about how to configure traffic interception for your WAAS network, see Chapter 5, "Configuring Traffic Interception."

Advantages and Disadvantages of Using Inline Interception

Inline interception requires using a WAE appliance that has the Cisco WAE Inline Network Adapter, Cisco Interface Module, or Cisco AppNav Controller Interface Module installed. In inline mode, the WAE can physically and transparently intercept traffic between the clients and the router. When using this mode, you physically position the WAE device in the path of the traffic that you want to optimize, typically between a switch and a router.

Because redirection of traffic is not necessary, inline interception simplifies deployment and avoids the complexity of configuring WCCP or PBR on the routers.

The inline adapter or module contains one or more pairs of LAN/WAN Ethernet ports each grouped into an inline or bridge group interface. If the inline adapter or module has multiple pairs of ports, it can connect to multiple routers if the network topology requires it.

The inline or bridge group interface transparently intercepts traffic flowing through it or bridges traffic that does not need to be optimized. It also uses a mechanical fail-safe design that automatically bridges traffic if a power, hardware, or unrecoverable software failure occurs.

Note AppNav Controller Interface Modules do not support automatic bypass mode to continue traffic flow in the event of a failure. For high availability, two or more AppNav Controller Interface Modules should be deployed in an AppNav cluster. For more information on using inline mode with the AppNav solution, see Chapter 4, "Configuring AppNav."

You can configure the inline or bridge group interface to accept traffic only from certain VLANs; for all other VLANs, traffic is bridged and not processed.

You can serially cluster WAE devices (not AppNav Controllers) in inline mode to provide higher availability in the event of a device failure. If the current optimizing device fails, the second WAE device in the cluster provides the optimization services. Deploying WAE devices in a serial inline cluster for the purposes of scaling or load balancing is not supported.

Any combination of traffic interception mechanisms on peer WAEs is supported. For example, you can use inline interception on the branch WAE and WCCP on the data center WAE. For complex data center deployments, we recommend that you use hardware-accelerated WCCP interception or load balancing with the Cisco Application Control Engine (ACE) and a WAAS AppNav deployment.

For more information on inline interception, see the "Using Inline Mode Interception" section on page 5-42.

Three elements can help ease traffic interception in data centers without using a WCCP-based approach:

- Multiple pairs of inline interfaces are available on certain WAE models:
  - WAVE-294/594/694/7541/7571/8541 models support one installed Cisco Interface Module, which can be configured with up to 16 inline ports in 8 inline groups, or one installed AppNav Controller Interface Module, which can be configured with up to 12 inline ports in 5 bridge groups.
  - WAE-674/7341/7371 models support dual inline Cisco WAE Inline Network Adapters, providing a total of 8 ports in 4 inline groups.
- Serial inline clustering of two WAEs (not AppNav Controllers) to support high availability.
- Interception ACLs to control what traffic is intercepted and what is passed through. For more information on interception ACLs, see the "Configuring Interception Access Control Lists" section on page 5-28.

Advantages and Disadvantages of Using WCCP-Based Routing

WCCP specifies interactions between one or more routers (or Layer 3 switches) and one or more application appliances, web caches, and caches of other application protocols. The purpose of the interaction is to establish and maintain the transparent redirection of selected types of traffic flowing through a group of routers. The selected traffic is redirected to a group of appliances.

WCCP allows you to transparently redirect client requests to a WAE for processing. The WAAS software supports transparent intercept of all TCP traffic.

To configure basic WCCP, you must enable the WCCP Version 2 service on the router and WAE or ANC in the data center and the router and WAE in the branch office. You do not need to configure all of the available WCCP features or services in order to get a WAE up and running.

Note You must configure the routers and WAEs to use WCCP Version 2 instead of WCCP Version 1 because WCCP Version 1 only supports web traffic (port 80).

WCCP is much simpler to configure than PBR. However, you need to have write access to the router in order to configure WCCP on the router, which typically resides in the data center and on the edge of the branch office. Another advantage of using WCCP is that you only need to perform a basic configuration of WCCP on your routers and WAEs in order to get your WAE up and running.

The WCCP Version 2 protocol also has a set of useful features built-in, for example, automatic failover and load balancing between multiple devices. The WCCP-enabled router monitors the liveliness of each WAE or ANC that is attached to it through the WCCP keepalive messages. If a WAE goes down, the

router stops redirecting packets to the WAE. When you use WCCP Version 2, the branch WAE is not made a single point of failure for the WAAS services. The router or ANC can also load balance the traffic among a number of branch WAEs.

You can use CLI commands to configure basic WCCP on both the routers and the WAEs, or you can use CLI commands to configure the router for WCCP and use the WAAS Central Manager GUI to configure basic WCCP on the WAEs.

We recommend that you use the WAAS CLI to complete the initial basic configuration of WCCP on your first branch WAE and data center WAE, as described in the *Cisco Wide Area Application Services Quick Configuration Guide*. After you have verified that WCCP transparent redirection is working properly, you can use the WAAS Central Manager GUI to centrally modify this basic WCCP configuration or configure additional WCCP settings (for example, load balancing) for a WAE (or group of WAEs). For more information, see the "Configuring WCCP on WAEs" section on page 5-11. After you have configured basic WCCP on the router, you can configure advanced WCCP features on the router, as described in the "Configuring Advanced WCCP Features on Routers" section on page 5-6.

Advantages and Disadvantages of Using PBR

PBR allows IT organizations to configure their network devices (a router or a Layer 4 to Layer 6 switch) to selectively route traffic to the next hop based on the classification of the traffic. WAAS administrators can use PBR to transparently integrate a WAE into their existing branch office network and data centers. PBR can be used to establish a route that goes through a WAE for some or all packets, based on the defined policies.

To configure PBR, you must create a route map and then apply the route map to the router interface on which you want the transparent traffic redirection to occur. Route maps reference access lists that contain explicit permit or deny criteria. The access lists define the traffic that is "interesting" to the WAE (that is, traffic that the network device should transparently intercept and redirect to the local WAE). Route maps define how the network device should handle "interesting" traffic (for example, send the packet to the next hop, which is the local WAE).

The following list summarizes the main advantages of using PBR instead of WCCP Version 2 to transparently redirect IP/TCP traffic to a WAE:

- PBR provides higher performance than WCCP Version 2 because there is no GRE overhead.
- By default PBR uses CEF when CEF is enabled on the router (PBR using CEF for fast switching of packets).
- PBR can be implemented on any Cisco IOS-capable router or switch that is running an appropriate version of the Cisco IOS software. We recommend that you use Cisco IOS Software Release 12.2 or later.
- PBR provides failover if multiple next-hop addresses are defined.

The following list summarizes the main disadvantages of using PBR instead of WCCP Version 2 to transparently redirect IP/TCP traffic to a WAE:

- PBR does not support load balancing between equal cost routes. Consequently, PBR does not provide scalability for the deployment location.
- PBR is more difficult to configure than WCCP Version 2. For an example of how to configure PBR for WAAS traffic, see the "Using Policy-Based Routing Interception" section on page 5-33.

Configuring WCCP or PBR Routing for WAAS Traffic

The primary function of WAAS is to accelerate WAN traffic. In general, WAAS accelerates TCP traffic. WAAS uses a symmetric approach for application optimization. A WAE that has application-specific and network-specific intelligence is placed on each side of the WAN. These WAEs are deployed out of the data path in both the branch office and the data center.

Traffic between the clients in the branch offices and the servers at the data center is transparently redirected through the WAEs based on a set of configured policies with no tunneling. The routers use WCCP Version 2 or PBR to transparently intercept and redirect traffic to the local WAE for optimization, redundancy elimination, and compression. For example, Edge-Router1 uses PBR or WCCP Version 2 to transparently redirect traffic to Edge-WAE1, the local WAE in the branch office. Core-Router1 uses PBR or WCCP Version 2 to transparently redirect traffic to the Core-WAE1, the local WAE in the data center.

Note In this sample deployment, the Edge-Router1 and Core-Router1 could be replaced with Layer 4 to Layer 7 switches, which are capable of redirecting traffic to the local WAE.

Figure 2-1 shows that the WAEs (Edge-WAE1 and Core-WAE1) must reside in an out-of-band network that is separate from the traffic's destination and source. For example, Edge-WAE1 is on a subnet separate from the clients (the traffic source), and Core-WAE1 is on a subnet separate from the file servers and application servers (the traffic destination). Additionally, you may need to use a tertiary interface (a separate physical interface) or a subinterface to attach a WAE to the router, which redirects traffic to it, to avoid an infinite routing loop between the WAE and the router. For more information on this topic, see the "Using Tertiary Interfaces or Subinterfaces to Connect WAEs to Routers" section on page 2-24.

Figure 2-1 Example of Using PBR or WCCP Version 2 for Transparent Redirection of All TCP Traffic to WAEs

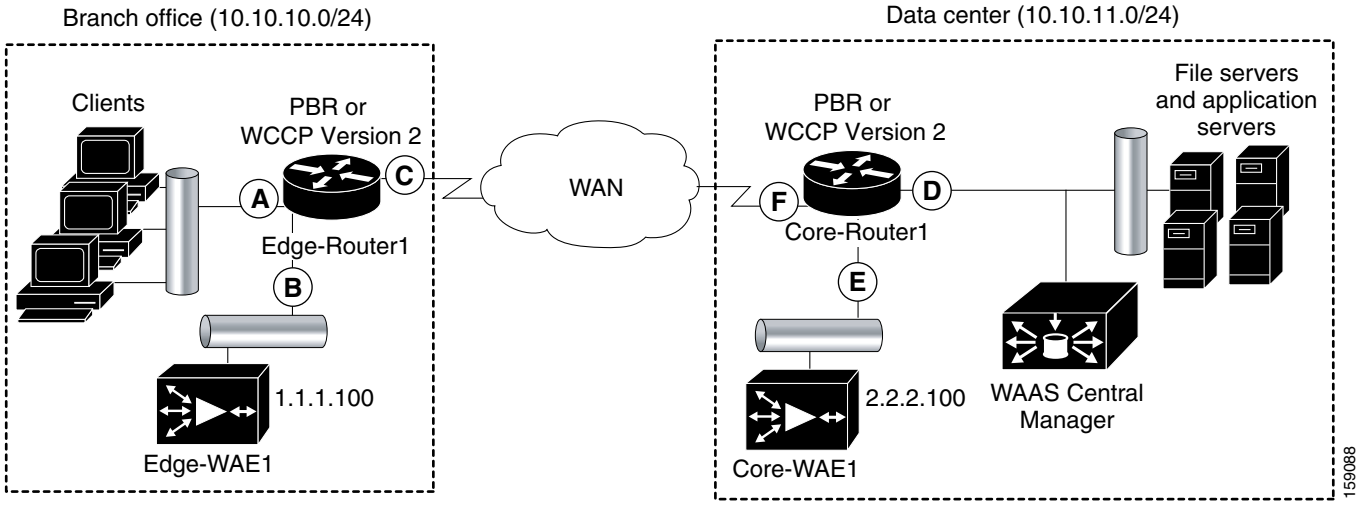

Table 2-1 provides a summary of the router interfaces that you must configure to use PBR or WCCP Version 2 to transparently redirect traffic to a WAE.

Table 2-1 Router Interfaces for WCCP or PBR Traffic Redirection to WAEs

Router interface	Description
Edge-Router1	
A	Edge LAN interface (ingress interface) that performs redirection on the outbound traffic.
B	Tertiary interface (separate physical interface) or a subinterface off of the LAN port on Edge-Router1. Used to attach Edge-WAE1 to Edge-Router1 in the branch office.
C	Edge WAN interface (egress interface) on Edge-Router1 that performs redirection on the inbound traffic.
Core-Router1	
D	Core LAN interface (ingress interface) that performs redirection on outbound traffic.
E	Tertiary interface or subinterface off of the LAN port on Core-Router1. Used to attach Core-WAE1 to Core-Router1 in the data center.
F	Core WAN interface (egress interface) on Core-Router1 that performs redirection on the inbound traffic.

This traffic redirection does not use tunneling; the full original quadruple (source IP address, source port number, destination IP address, and destination port number) of the TCP traffic is preserved end to end. The original payload of the TCP traffic is not preserved end to end because the primary function of WAAS is to accelerate WAN traffic by reducing the data that is transferred across the WAN. This change in payload can potentially impact features on the router (which is performing the WCCP or PBR redirection) that needs to see the actual payload to perform its operation (for example, NBAR). For more information on this topic, see the "WAAS and Cisco IOS Interoperability" section on page 2-11.

Using WCCP or PBR at both ends with no tunneling requires that traffic is intercepted and redirected not only in the near-end router but also at the far-end router, which requires four interception points as opposed to two interception points in a tunnel-based mode.

You can enable packet redirection on either an outbound interface or inbound interface of a WCCP-enabled router. The terms *outbound* and *inbound* are defined from the perspective of the interface. Inbound redirection specifies that traffic should be redirected as it is being received on a given interface. Outbound redirection specifies that traffic should be redirected as it is leaving a given interface.

If you are deploying WAN optimization in your WAAS network, then you must configure the router and WAE for WCCP Version 2 and the TCP promiscuous mode service (WCCP Version 2 services 61 and 62 by default).

Note Services 61 and 62 are always enabled together when configuring TCP promiscuous on the WAE. Services 61 and 62 must be defined and configured separately when configuring TCP promiscuous on the network device (router, switch, or other). Service 61 distributes traffic by source IP address, and service 62 distributes traffic by destination IP address. The service IDs are configurable; 61 and 62 are the defaults.

The TCP promiscuous mode service intercepts all TCP traffic that is destined for any TCP port and transparently redirects it to the WAE. The WCCP-enabled router uses service IDs 61 and 62 to access this service. The service IDs used on the router must match those on the WAE if different service IDs than the defaults are configured.

By default, the IP Protocol 6 is specified for the TCP promiscuous mode service. Consequently, the routers that have been configured to the TCP promiscuous mode service will intercept and redirect all TCP traffic destined for any TCP port to the local WAE. Because the TCP promiscuous mode service is configured on the WAE, the WAE will accept all of the TCP traffic that is transparently redirected to it by specified WCCP routers (for example, Edge-WAE1 will accept all TCP traffic that Edge-Router1 redirects to it). In the branch office, you can intercept packets at the edge LAN and WAN interfaces on the edge routers and redirect the TCP traffic to the local WAE (the branch WAE). In the data center, you can intercept packets at the core LAN and WAN interfaces on the core routers and redirect the TCP traffic to the local WAE (the data center WAE). For more information, see the "Configuring WAEs as Promiscuous TCP Devices in a WAAS Network" section on page 2-24.

Configure packet redirection on inbound interfaces of branch software routers whenever possible. Inbound traffic can be configured to use Cisco Express Forwarding (CEF), distributed Cisco Express Forwarding (dCEF), fast forwarding, or process forwarding.

Note CEF is required for WCCP and must be enabled on the router.

To enable packet redirection on a router's outbound or inbound interface using WCCP, use the **ip wccp redirect** interface configuration command.

Caution The **ip wccp redirect** interface command has the potential to affect the **ip wccp redirect exclude in** command. If you have **ip wccp redirect exclude in** set on an interface and you subsequently configure the **ip wccp redirect in** command, the **exclude in** command is overridden. If you configure the **exclude in** command, the **redirect in** command is overridden.

This section contains the following topics:

Configuring WAEs as Promiscuous TCP Devices in a WAAS Network

In order for the WAE to function as a promiscuous TCP device for the TCP traffic that is transparently redirected to it by the specified WCCP Version 2 routers, the WAE uses WCCP Version 2 services 61 and 62 by default, though the service IDs are configurable. The WCCP services are represented by the canonical name tcp-promiscuous on the WAE CLI and TCP Promiscuous in the WAAS Central Manager GUI. (See Figure 5-3.)

For instructions on how to perform a basic WCCP configuration for a WAAS network, see the *Cisco Wide Area Application Services Quick Configuration Guide*. For instructions about how to use the WAAS Central Manager GUI to modify the basic WCCP configuration for a WAE, see the "Configuring WCCP on WAEs" section on page 5-11.

Using Tertiary Interfaces or Subinterfaces to Connect WAEs to Routers

If you plan to use WCCP Version 2 or PBR to transparently redirect TCP traffic to a WAE, make sure that the WAE is not attached to the same segment as the router interface on which the traffic redirection is to occur. Otherwise, an infinite routing loop between the router and the WAE will occur. These infinite routing loops occur because there is no way to notify the router to bypass the interception and redirection after it has redirected the traffic to the WAE the first time; the router will continuously redirect the same intercepted traffic to the local WAE, creating the infinite routing loop.

Note The WCCP GRE return and generic GRE egress methods allow you to place WAEs on the same VLAN or subnet as clients and servers. For information on configuring these egress methods, see the "Configuring Egress Methods for WCCP Intercepted Connections" section on page 5-29.

For example, if you attach Edge-WAE 1 to the same segment (subnet) as the LAN router interface on which the PBR or WCCP traffic redirection occurs in the branch office, there will be an infinite routing loop between Edge-Router1 and Edge-WAE1. If you attach Core-WAE1 to the same segment (subnet) as the LAN router interface on which the PBR or WCCP traffic redirection occurs in the data center, there will be an infinite routing loop between Core-Router1 and Core-WAE1.

To avoid an infinite routing loop between the router and its local WAE, connect the WAE to the router through a tertiary interface (a separate physical interface) or a subinterface (a different virtual subinterface) from the router's LAN port. By using a tertiary interface or a subinterface to connect a WAE to the router that is performing the PBR or WCCP redirection, the WAE has its own separate processing path that has no Cisco IOS features enabled on it. In addition, this approach simplifies the process of integrating WAEs into an existing network. Because the WAEs are being connected to the routers through a tertiary interface or subinterface that has no Cisco IOS features enabled on it, the Cisco IOS features that are already enabled on your existing Cisco-enabled network elements (for example, Edge-Router1 or Core-Router1) will generally not be affected when you connect WAEs to these routers. For more information about WAAS and Cisco IOS interoperability, see the "WAAS and Cisco IOS Interoperability" section on page 2-11.

See the *Cisco Wide Area Application Services Quick Configuration Guide* for an example of how to use a subinterface to properly attach a local WAE to the router that is redirecting TCP traffic to it.

Access Lists on Routers and WAEs

You can optionally configure the router to redirect traffic from your WAE based on access lists that you define on the router. These access lists are also referred to as redirect lists. For information about how to configure access lists on routers that will be configured to transparently redirect traffic to a WAE, see the "Configuring IP Access Lists on a Router" section on page 5-9.

Note IP access lists on routers have the highest priority followed by IP ACLs that are defined on the WAEs, and then interception ACLs that are defined on the WAEs.

This section contains the following topics:

IP ACLs on WAEs

In a centrally managed WAAS network environment, administrators need to be able to prevent unauthorized access to various devices and services. The WAAS software supports standard and extended IP access control lists (ACLs) that allow you to restrict access to or through particular interfaces on a WAAS device. For more information, see Chapter 9, "Creating and Managing IP Access Control Lists for WAAS Devices."

Note IP ACLs that are applied on interfaces, and WCCP ACLs, always take precedence over any interception ACLs and WAAS application definitions that have been defined on the WAE.

Interception ACLs on WAEs

You can configure an interception ACL to control what incoming traffic across all interfaces is to be intercepted by a WAE device. Packets that are permitted by the ACL are intercepted by the WAE and packets that are denied by the ACL are passed through the WAE without processing. By configuring interception ACLs on the WAE, you can control traffic interception without modifying the router configuration.

An interception ACL can be used both with WCCP and inline interception.

Interception ACLs that are defined on a WAE always take precedence over any WAAS application definitions that have been defined on the WAE, but they are applied after interface ACLs and WCCP ACLs.

For information about how to configure an interception ACL for a WAE, see the "Configuring Interception Access Control Lists" section on page 5-28.

WAAS Login Authentication and Authorization

In the WAAS network, administrative login authentication and authorization are used to control login requests from administrators who want to access a WAAS device for configuring, monitoring, or troubleshooting purposes.

Login authentication is the process by which WAAS devices verify whether the administrator who is attempting to log in to the device has a valid username and password. The administrator who is logging in must have a user account registered with the device. User account information serves to authorize the user for administrative login and configuration privileges. The user account information is stored in an AAA database, and the WAAS devices must be configured to access the particular authentication server (or servers) where the AAA database is located. When the user attempts to log in to a device, the device compares the person's username, password, and privilege level to the user account information that is stored in the database.

The WAAS software provides the following authentication, authorization, and accounting (AAA) support for users who have external access servers (for example, RADIUS, TACACS+, or Windows domain servers), and for users who need a local access database with AAA features:

- *Authentication* (or *login authentication*) is the action of determining who the user is. It checks the username and password.
- *Authorization* (or *configuration*) is the action of determining what a user is allowed to do. It permits or denies privileges for authenticated users in the network. Generally, authentication precedes authorization. Both authentication and authorization are required for a user log in.
- *Accounting* is the action of keeping track of administrative user activities for system accounting purposes. In the WAAS software, AAA accounting through TACACS+ is supported.

For more information, see the "Configuring AAA Accounting for WAAS Devices" section on page 7-31.

WAAS Administrator Accounts

In a centrally managed WAAS network, administrator accounts can be created for access to the WAAS Central Manager and, independently, for access to the WAEs that are registered with the WAAS Central Manager. There are two distinct types of accounts for WAAS administrators:

- Role-based accounts—Allows users to access the WAAS Central Manager GUI, the WAAS Central Manager CLI, and the WAE Device Manager GUI. The WAAS software has a default WAAS system user account (username is admin and password is default) that is assigned the role of administrator.
- Device-based CLI accounts—Allow users to access the WAAS CLI on a WAAS device. These accounts are also referred to as local user accounts.

Note An administrator can log in to the WAAS Central Manager device through the console port or the WAAS Central Manager GUI. An administrator can log in to a WAAS device that is functioning as a data center or branch WAE through the console port or the WAE Device Manager GUI.

A WAAS device that is running WAAS software comes with a predefined superuser account that can be used initially to access the device. When the system administrator logs in to a WAAS device before authentication and authorization have been configured, the administrator can access the WAAS device by using the predefined superuser account (the predefined username is admin and the predefined password is default). When you log in to a WAAS device using this predefined superuser account, you are granted access to all the WAAS services and entities in the WAAS system.

After you have initially configured your WAAS devices, we strongly recommend that you immediately change the password for the predefined superuser account (the predefined username is *admin*, the password is *default*, and the privilege level is superuser, privilege level 15) on each WAAS device. For instructions on how to use the WAAS Central Manager GUI to change the password, see the "Changing the Password for Your Own Account" section on page 8-6.

Logically Grouping Your WAEs

To streamline the configuration and maintenance of WAEs that are registered with a WAAS Central Manager, you can create a logical group and then assign one or more of your WAEs to the group. Groups not only save you time when configuring multiple WAEs, but they also ensure that configuration settings are applied consistently across your WAAS network. For example, you can set up a WinAuth group that defines the standard Windows authentication configuration that is wanted for all of the WAEs in that group. After you define the WinAuth settings once, you can centrally apply those values to all of the WAEs in the WinAuth group instead of defining these same settings individually on each WAE.

With the WAAS Central Manager GUI, you can easily organize your branch and data center WAEs into device groups, which are a collection of WAEs that share common qualities and capabilities. Setting up groups based on their authentication settings is an example of a device group.

When you create a device group, you need to identify the unique characteristics that distinguish that group of WAEs from others in your network. For example, in larger WAAS deployments one set of WAEs may need to be configured with authentication settings that are different from another set of WAEs in your WAAS network. In this case, you would create two device groups that each contain different authentication settings, and then assign your WAEs to the most appropriate group.

If you have WAEs that reside in different time zones, you can also create device groups based on geographic regions so that the WAEs in one group can have a different time zone setting from the WAEs in another group.

In smaller WAAS deployments where all WAEs can be configured with the same settings, you may only need to create one general device group. This practice allows you to configure settings for the group, then apply those settings consistently across all your WAEs.

Note The AllWAASGroup and AllWAASExpressGroup are default device groups that automatically contain all WAAS and WAAS Express devices. In these or any other device groups, you should configure only the settings that you want to be consistent across all the devices in the group. Settings that apply to a single device should be configured on that device only and not on the device group.

By default, WAAS Central Manager allows you to assign a device to multiple device groups. Before you create a device group, make sure you understand the unique properties that you want the group to contain.

WAAS Central Manager allows you to create locations that you can associate with a WAAS device. You assign a device to a location when you first activate the device. The main purpose of assigning a WAAS device to a location is to help you identify a WAAS device by the physical region in which it resides. Locations are different from device groups because devices do not inherit settings from locations.

You assign a device to a location when you activate the device as described in the *Cisco Wide Area Application Services Quick Configuration Guide*. For more information about logically grouping your WAEs, see Chapter 3, "Using Device Groups and Device Locations."

Data Migration Process

If you have an existing network, there are some steps to take before setting up your WAAS network. The first step in the data migration process is to back up the data at the branch offices and restore it to the data center.

After you back up data to the data center, you preload the cache (called *preposition*) with the files for which you want to provide the fastest access. Set up the files from your branch office file server to the WAEs that are also located in the same branch office. You can then remove the file servers from the branch offices and point to the data center file server.

The final step in the data migration process is to set the CIFS policies.

When doing the data migration process, note the following restrictions:

- Prepositioning only works in a CIFS environment with the CIFS accelerator (it is not supported by the SMB accelerator).
- The topology for the file server at the data center must be identical to the topology that existed on the branch file server.
- Resource credentials (such as ACLs) are not automatically migrated. Two options are available:
  - You can use backup or restore software to restore an initial backup of the tree to the target server. This practice allows both the creation of ACLs as well as the creation of the initial file set that Rsync can take as an input for diff calculations. The replication inherits existing ACLs in that tree.
  - The other option is to perform a first run of Robocopy (including data and permissions), and then continue with sync iterations using Rsync.

  After replicating, use one of Microsoft's tools for copying only ACLs (no data) onto the replicated tree. You can use Robocopy.exe for copying directory tree or file ACLs and Permcopy.exe to copy share permissions.

- The migration size must be less than the cache size of the branch WAE.

CHAPTER 3

Using Device Groups and Device Locations

This chapter describes the types of device groups supported by the WAAS software and how to create groups that make it easier to manage and configure multiple devices at the same time. This chapter also discusses how to use device locations.

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE and WAVE appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

This chapter contains the following sections:

About Device Groups

When you create a device group, you need to identify the unique characteristics that distinguish that group of devices from others in your network. For example, in larger WAAS deployments, one set of devices may need to be configured with authentication settings that are different from another set of devices in your WAAS network. In this situation, you would create two device groups that each contain different authentication settings, and then assign your devices to the most appropriate group.

If you have devices that reside in different time zones, you can also create device groups based on geographic regions so that the devices in one group can have a different time zone setting from the devices in another group.

In smaller WAAS deployments where all devices can be configured with the same settings, you may only need to create one general device group. This setup allows you to configure settings for the group, and then apply those settings consistently across all your WAAS devices.

Groups not only save you time when configuring multiple devices, but they also ensure that configuration settings are applied consistently across your WAAS network.

There are two types of device groups: WAAS Device Groups and WAAS Express Device Groups. These groups are explained in more detail in the "Creating a New Device Group" section on page 3-3.

When you register a WAAS device with the WAAS Central Manager, that device automatically joins the AllWAASGroup, which is the default device group on the system for WAAS devices. If you create additional device groups, you need to decide if you want your devices to belong to more than one group

(the default AllWAASGroup and the new device group you create). If you only want a device to belong to a device group that you create, make sure that you remove the device from the default AllWAASGroup. WAAS Express devices automatically join the default AllWAASExpressGroup device group when they are registered with the Central Manager.

WAAS devices and WAAS Express devices cannot be mixed in the same device group. You choose the device group type when you create the group and it cannot be changed. When you create a WAAS Express type of device group, you can copy policies from an existing WAAS or WAAS Express group, but policies cannot be copied after creation.

Working with Device Groups

This section contains the following topics:

Creating a Device Group

This section contains the following topics:

Table 3-1 describes the process for creating a new device group.

Table 3-1 Checklist for Creating a Device Group

Task	Additional Information and Instructions
1. Create a new device group.	Defines general information about the new group, such as the group name, group type, and whether all newly activated devices are assigned to this group. For more information, see the "Creating a New Device Group" section on page 3-3.
2. Configure the settings of the new device group.	Specifies the settings that are unique to this device group. All devices that are a member of this group will automatically inherit these settings. For more information, see the "Configuring the Settings for a Device Group" section on page 3-4.

Table 3-1 Checklist for Creating a Device Group (continued)

Task	Additional Information and Instructions
3. Assign devices to the device group.	Assigns devices to the group so they can inherit the group settings. For more information, see the "Assigning Devices to a Configuration Device Group" section on page 3-5.

Creating a New Device Group

Before you create a device group, make sure you understand the unique properties that you want the group to contain. For example, you may want to set up two device groups that have different authentication settings or different time zone settings.

To create a device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Device Groups > All Device Groups**. The Device Groups window appears.

From this window you can perform the following tasks:

- Click the **Edit** icon next to the device group that you want to modify.
- Create a new device group as described in the steps that follow.

Step 2 Click the **Create New Device Group** icon in the taskbar. The Creating New Device Group window appears.

Step 3 In the Name field, enter the name of the device group.

The name must be unique and should be a name that is useful in distinguishing the device group from others on your system. The name cannot contain characters other than letters, numbers, period, hyphen, underscore, and space.

Step 4 Choose either WAAS or WAAS Express for the Configuration Group Type. This sets the type of devices that the group can contain. A WAAS Express group can contain only WAAS Express devices. A WAAS group can contain all types of devices except for WAAS Express devices.

Step 5 Check the **Automatically assign all newly activated devices to this group** check box to set this device group as the default device group for all newly activated devices.

Step 6 If you chose the WAAS Express group type, you can copy policies from another existing group by choosing the group in the Copy Policies from the device group drop-down list (only shown when creating a WAAS Express group). If you copy policies from a WAAS group, only basic optimization policies are copied, not application acceleration policies.

Step 7 (Optional) Enter comments about the group in the Comments field. The comments that you enter will appear in the Device Group window.

Step 8 Click **Submit**.

The page refreshes with additional options.

Note The Pages configured for this device group arrow lists the configuration windows in the WAAS Central Manager GUI that have been configured for this device group. Because this is a new device group, no pages will appear in this list.

Step 9 (Optional) Customize the menu options for this device group by completing the following steps. Use this feature to remove from view any configuration windows that you do not need for that particular device group:

a. Click the **Select pages to hide from table of contents for this device group** arrow.

   A list of windows in the WAAS Central Manager GUI appears.

b. Check the windows that you want to hide for this device group. You can click the folder icon next to a window to display its child windows.

c. Click **Submit**.

Step 10 Configure the settings for this device group as described in the "Configuring the Settings for a Device Group" section.

Configuring the Settings for a Device Group

After creating a device group, you need to configure the settings that you want to be unique to this group.

If you have a general device group that contains all your WAAS devices of a specific type, configure only the settings that you want to be consistent across all the devices of that type. Settings that apply to a single device should be configured on that device only and not on the device group.

To configure settings for a device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Device Groups >** *device-group-name*.

The Modifying Device Group window appears.

Step 2 Click the **Pages configured for this device group** arrow button to view which configuration windows have already been configured for the group.

A list of pages that are configured for that device group appears. If this is a new device group or if there are no pages configured for this device group, the list displays Null.

Step 3 Customize the menu options for this device group by completing the following steps:

a. Click the **Select pages to hide from table of contents of this device group** arrow.

   A list of windows in the WAAS Central Manager GUI appears.

b. Place a check next to the windows that you want to hide for this device group. Use this feature to remove from view any configuration windows that you do not need for this particular device group.

Step 4 Use the menu bar to choose each configuration option that you want to modify for this device group.

If the configuration option has not been configured for this device group, the message "There are currently no settings for this group" appears at the top of the window.

Step 5 Make the necessary changes on the configuration option window, and click **Submit** when finished.

After a particular setting is configured, the configuration window is listed under Pages configured for this device group in the Modifying Device Group window.

Step 6 Assign devices to this new group as described in the "Assigning Devices to a Configuration Device Group" section on page 3-5.

Assigning Devices to a Configuration Device Group

After you create a configuration device group, you need to assign devices to the group. The WAAS Central Manager GUI provides two methods to assign devices to a configuration group. You can either select the device first, then assign a group to the device, or you can select the device group first, then assign devices to the group.

The procedures in this section describe how to assign devices to a group. To assign a group to a device, choose **Devices >** *device-name* and choose **Assign Device Groups** from the device-name menu. You can then assign a group to the device using the same method described in steps 4 and 5 below.

You cannot assign the WAAS Central Manager to a device group. You must configure the WAAS Central Manager separately from other devices.

You cannot assign WAAS Express devices to a WAAS group and you cannot assign WAAS devices to a WAAS Express group. Invalid devices are not shown in the device list when assigning devices to groups.

Note By default, all devices automatically join either the AllWAASGroup or AllWAASExpressGroup when they are activated. If you do not want a device to belong to two different device groups, you should unassign the device from the All...Group before you assign the device to a custom device group.

Use care when you are assigning devices that have different WAAS software versions to a device group. Some features configured for a device group may not be supported by all devices in the group or, in some cases, devices may be prevented from joining the group if the group is configured with policies that they cannot support. In such cases, we recommend that you upgrade all devices to the same software version or create different device groups for devices with incompatible versions.

To assign a device to a device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Device Groups >** *device-group-name*.

The Modifying Device Group window appears.

Step 2 Choose *device-group-name* **> Assign Devices**.

The WAE/WAAS Express Assignments window appears, displaying the devices assigned to various locations. If you are editing a WAAS group, only WAAS devices are shown. If you are editing a WAAS Express group, only WAAS Express devices are shown.

The assignments window lets you filter your view of the items in the list. Filtering allows you to find items in the list that match the criteria that you set.

Step 3 Assign a device to the device group by doing either of the following:

- Click in the taskbar to assign all available devices to the group.
- Click next to each device that you want to assign to the group. The icon changes to when selected.

Step 4 Click **Submit**.

A green check mark appears next to the assigned devices.

Step 5 Click the **Unassign** icon (green check mark) next to the name of the device that you want to remove from the device group. Alternatively, you can click the **Remove all** icon in the taskbar to remove all devices from the selected device group. Click **Submit**.

Deleting a Device Group

To delete a device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Device Groups >** *device-group-name*. The Modifying Device Group window appears.

Step 2 In the taskbar, click the **Delete Device Group** icon. You are prompted to confirm your decision to delete the device group.

Step 3 To confirm your decision, click **OK**.

Viewing Device Group Assignments

The WAAS Central Manager GUI allows you to view the groups that a device belongs to, as well as the devices that belong to a specific group. This section describes both of these procedures.

To view the groups that a device belongs to, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

The Device Dashboard window appears.

Step 2 In the Assignments field on the Device Dashboard window, click the link that displays the groups to which the device is assigned.

The Device Group Assignments page appears, which shows all the device groups in your WAAS network that match the device type (WAAS or WAAS Express). The device is assigned to the device groups with a green check mark next to them.

You can also go to the Device Group Assignments window by choosing the Assign Device Groups option in the menu bar.

To view the devices that are assigned to a specific group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Device Groups >** *device-group-name*.

The Modifying Device Group window appears.

Step 2 Choose *device-group-name* **> Assign Devices**.

The WAE/WAAS Express Assignments window appears, which shows all the WAAS or WAAS Express devices on your WAAS network. The devices with a green check mark next to them are assigned to this group.

Viewing the Device Groups List

The Device Groups window lists all the device groups that have been created in your WAAS network. To view this list, choose **Device Groups > All Device Groups** in the WAAS Central Manager menu bar.

This window displays the following information about each device group:

- Type of device group (WAAS Configuration Group or WAAS Express Configuration Group).
- Any comments that were entered when the device group was created.

From this window, you can perform the following tasks:

- Create a new device group. For more information, see the "Creating a New Device Group" section on page 3-3.
- Modify the settings of a device group by clicking the **Edit** icon next to the group that you want to edit.

Enabling or Disabling Device Group Overlap

By default, you can assign a device to multiple device groups. You can disable this functionality so a device can only belong to one device group, which eliminates the possibility of a device inheriting settings from more than one group.

To enable or disable device group overlap, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Configure > Global > System Properties**.

The Config Properties window appears.

Step 2 Click the **Edit** icon next to the property name DeviceGroup.overlap.

The Modifying Config Property, DeviceGroup.overlap window appears.

Step 3 From the Value drop-down list, choose either **true** or **false**. (The default is true.)

When you disable device group overlap (set to false), existing overlapping device groups are retained and continue to be handled as though overlap were enabled; however, any newly added groups do not allow overlapping, and new devices cannot be added to the existing overlapping groups.

Step 4 Click **Submit**.

Overriding Group Configuration Settings

The WAAS Central Manager GUI provides the following methods to override the current group configuration on a device:

- Forcing Device Group Settings on All Devices in the Group, page 3-7
- Selecting Device Group Precedence, page 3-8
- Overriding the Device Group Settings on a Device, page 3-8

Forcing Device Group Settings on All Devices in the Group

To force a device group configuration across all devices in the group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Device Groups >** *device-group-name*.

The Modifying Device Group window appears.

Step 2 Click the **Force Group Settings** icon in the taskbar.

The WAAS Central Manager GUI displays the following message:

```
The action will apply all settings configured for this device group to all the WAEs/WAAS
Express assigned to it. Do you wish to continue?
```

Step 3 To force group settings across all devices in the device group, click **OK**.

Step 4 Click **Submit**.

Selecting Device Group Precedence

When a device belongs to multiple device groups that have conflicting settings, the device automatically inherits the settings from the device group that was most recently changed. For a more detailed description of how a device inherits settings when it belongs to multiple device groups, see the "Understanding the Impact of Assigning a Device to Multiple Device Groups" section on page 3-9.

When a configuration conflict occurs, you can edit a device's configuration on a page-by-page basis and select which device group's settings should take precedence.

To select the device group precedence, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

The Device Dashboard window appears.

Step 2 From the menu bar, choose the configuration option that contains the conflicting settings.

A drop-down list appears in the taskbar at the top of the window. This drop-down list allows you to select the device group that you want this configuration window to inherit settings from. The device group that is currently selected is the device group that has precedence.

Step 3 From the drop-down list, choose the device group that you want this configuration page to inherit settings from, and click **Submit**.

The configuration window changes to reflect the settings associated with the selected device group.

Overriding the Device Group Settings on a Device

The WAAS Central Manager GUI allows you to override the device group settings and specify new settings that are unique to that device.

To override the device group settings on a device, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

The Device Dashboard window appears.

Step 2 From the menu bar, choose the configuration option that contains the device group settings you want to override.

Step 3 Click the **Override Group Settings** icon in the taskbar.

The settings in the configuration window are enabled.

Note The Override Group Settings icon only appears on configuration windows that have been modified on the associated device group.

Step 4 Make the necessary changes to the configuration window, and click **Submit**.

The device is now configured with settings that are different from the device group it belongs to.

Note The Force Settings on all Devices in Group icon appears in the device group view of an overridden configuration window. You can click this icon to reapply the device group settings to all devices in the device group.

Step 5 To reapply the device groups settings to this configuration window, choose the device group from the drop-down list in the taskbar, and click **Submit**.

Understanding the Impact of Assigning a Device to Multiple Device Groups

If a device belongs to multiple device groups, a configuration conflict might occur if the groups are not configured exactly the same. In this case, the device will inherit the settings from the device group that was most recently changed. In some cases, however, a device can retain settings from more than one device group depending on how the changes were implemented.

The following scenario describes how a device can retain settings from multiple device groups:

Action 1: Device A is assigned to Device Group 1 (DG1).

Result: Device A automatically inherits all the configuration settings of DG1.

Action 2: Device A is assigned to Device Group 2 (DG2) so it now belongs to two device groups (DG1 and DG2).

Result: Device A inherits all the settings from DG2, but it remains a member of DG1.

Action 3: The standard time zone setting on DG1 is changed to America New York.

Result: The time zone of Device A changes to America New York, but the device maintains all its other configuration settings from DG2.

In this scenario, Device A's configuration is a hybrid of DG1 and DG2. If you want to specify which device group settings a device should inherit, you can use the override features described in the "Overriding Group Configuration Settings" section on page 3-7.

Working with Device Locations

The WAAS Central Manager GUI allows you to create locations that you can associate with a WAAS device. You assign a device to a location when you first activate the device. The main purpose of assigning a device to a location is to help you identify a WAAS device by the physical region in which it resides. Locations are different from device groups because devices do not inherit settings from the location to which they belong.

You can view reports that aggregate data from all the devices in a particular location. For more information, see the "Location Level Reports" section on page 17-36.

You assign a device to a location when you activate the device as described in the "Modifying Device Properties" section on page 10-1.

You can work with locations by performing these tasks:

- Creating Locations, page 3-10
- Deleting Locations, page 3-10
- Viewing the Location Tree, page 3-11

Creating Locations

To create a new location or modify an existing one, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Locations > All Locations**. The Locations window appears.

Step 2 In the taskbar, click the **Create New Location** icon.

The Creating New Location window appears.

Step 3 In the Name field, enter a location name.

The name can contain letters, numbers, period, hyphen, underscore, and space.

Step 4 From the Parent Location drop-down list, choose a parent location (or choose **None**).

A location with no parent is a level 1 location. A location with a level 1 parent becomes a level 2 location, and so forth. The location level is displayed after you choose a parent location (or choose **None**) and click **Submit** to save the configuration.

Step 5 (Optional) In the Comments field, enter comments about the location.

Step 6 Click **Submit**.

Step 7 Modify a location by going to the Locations window and clicking the **Edit** icon next to the name of the location that you want to modify.

Step 8 Assign a device to this location. For more information, see the "Modifying Device Properties" section on page 10-1.

Deleting Locations

You can delete locations as needed, as long as they are not the root locations of activated WAAS devices.

Note If a location has a device assigned to it, you can first assign the device to another location and then delete the original location.

To delete a location, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Locations >** *location-name*.

The Modifying Location window appears.

Step 2 In the taskbar, click the **Delete Location** icon. You are asked to confirm your decision to delete the location.

Step 3 To confirm the action, click **OK**. The location is deleted.

Viewing the Location Tree

The location tree represents the network topology you configured when you assigned a parent to each location. The WAAS Central Manager GUI graphically displays the relationships between the locations configured in your WAAS network.

To view the location tree, choose **Locations > All Locations**. In the taskbar, click the **Location Trees** button.

CHAPTER 4

Configuring AppNav

This chapter describes how to configure Cisco WAAS AppNav, which is a hardware and software solution that simplifies network integration of WAN optimization and overcomes challenges with provisioning, visibility, scalability, asymmetry, and high availability.

This chapter includes the following topics:

Information About AppNav

AppNav greatly reduces dependency on the intercepting switch or router by distributing traffic among WAAS devices for optimization using a powerful class and policy mechanism. You can use WAAS nodes (WNs) to optimize traffic based on sites and/or applications.

The AppNav solution has the ability to scale up to available capacity by taking into account WAAS device utilization as it distributes traffic among nodes. Also, the solution provides for high availability of optimization capacity by monitoring node overload and liveliness and by providing configurable failure and overload policies.

This section includes the following sections:

System Components

The AppNav solution consists of the following components (see Figure 4-1):

- AppNav Controller (ANC)—A WAAS appliance with a Cisco AppNav Controller Interface Module that intercepts network traffic and, based on an AppNav policy, distributes that traffic to one or more WAAS nodes (WNs) for optimization.

- AppNav Controller Group (ANCG)—A group of AppNav Controllers within one service context that together provide the necessary intelligence for handling asymmetric flows and providing high availability. The ANCG is configured on the ANC. An ANCG can have up to eight ANCs.
- WAAS Node (WN)—A WAAS optimization engine (WAE or WAVE appliance, NME-WAE or SM-SRE network module, or vWAAS instance, but not a WAAS Express device) that optimizes and accelerates traffic according to the optimization policies configured on the device. You can have up to 32 WNs in the service context. (In the CLI, a WAAS node is also known as a service node.)
- WAAS Node Group (WNG)—A group of WAAS nodes within a service context that services a particular set of traffic flows identified by AppNav policies. The WNG is configured on the ANC. You can have up to 32 WNGs in the service context. (In the CLI, a WAAS node group is also known as a service node group.)
- AppNav Cluster—The group of all ANC and WN devices within a service context.
- Service Context—The topmost entity that groups together one AppNav Controller Group (ANCG), one or more WAAS node groups (WNGs), and an associated AppNav policy. The service context is configured on the ANC.

Figure 4-1 AppNav Solution Components

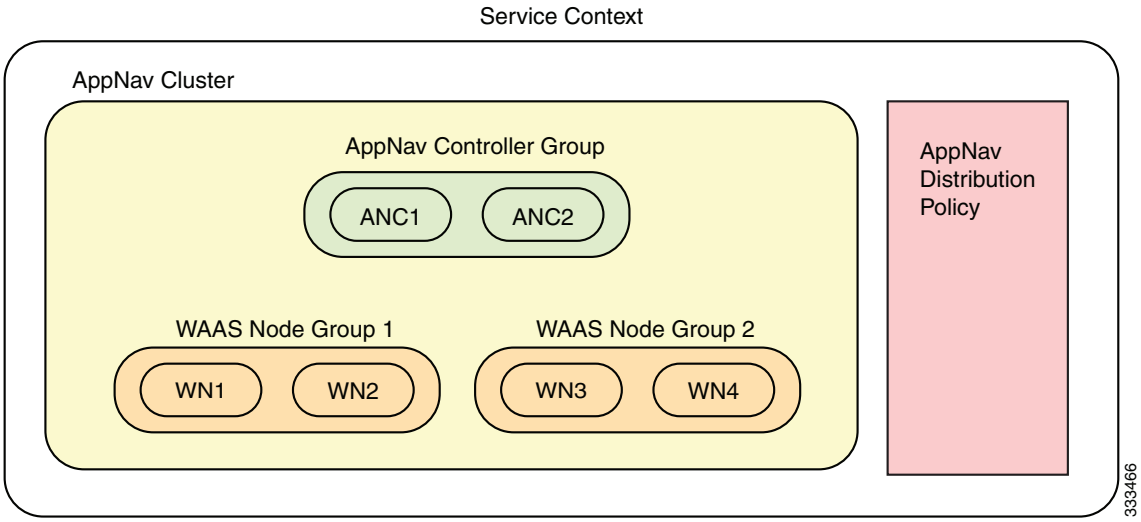

Within a service context, WAAS devices can operate in one of two modes:

- Application accelerator—The device serves only as a WN within the service context. It receives traffic from the ANC, optimizes the traffic, and returns the traffic to the ANC to be delivered to its destination. The WN can be any kind of WAAS device or vWAAS instance.
- AppNav Controller—The device operates as an ANC that intercepts network traffic and, based on a flow policy, distributes that traffic to one or more WNs for optimization. Only a WAVE appliance that contains a Cisco AppNav Controller Interface Module can operate as an ANC. An ANC can also operate as a WN and optimize traffic as part of a WNG.

AppNav Controller Deployment Models

You can deploy AppNav Controllers in your network in two ways (see Figure 4-2):

- In-path—The ANC is physically placed between one or more network elements, enabling traffic to traverse a bridge group configured on the device in inline mode.
- Off-path—The ANC works with the network infrastructure to intercept traffic through the Web Cache Communication Protocol (WCCP).

The ANC provides the same features in both in-path and off-path deployments. In either case, only ANCs participate in interception from the switch or router. The ANCs then distribute flows to WNs using a consistent and predictable algorithm that considers configured policies and WN utilization.

In Figure 4-2, WAAS Nodes could be attached to either or both switches in the diagrams.

Figure 4-2 Deployment Models

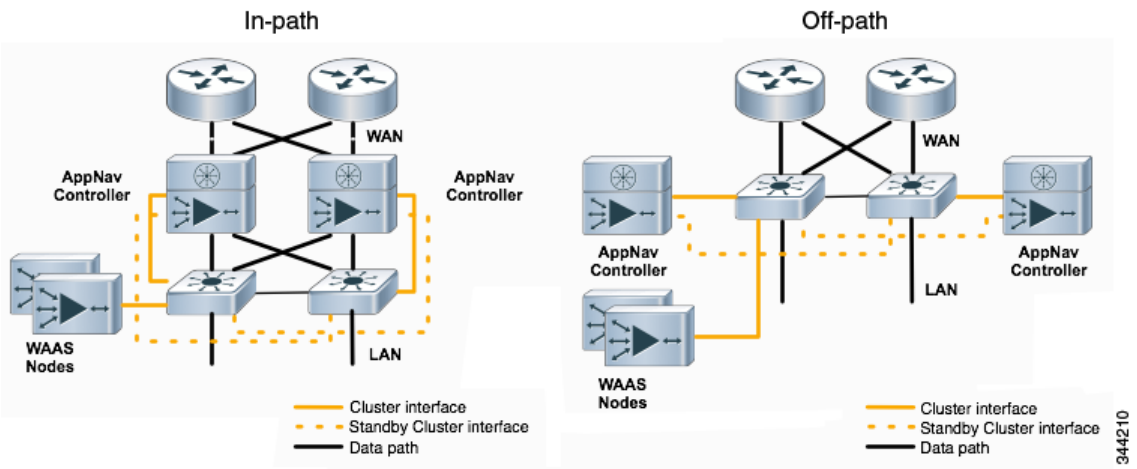

AppNav Controller Interface Modules

A WAAS appliance operating as an ANC requires a Cisco AppNav Controller Interface Module, which is similar to a standard WAVE appliance interface module but contains additional hardware, including a network processor and high speed ternary content addressable memory (TCAM), to provide intelligent and accelerated flow handling. The following AppNav Controller Interface Modules are supported:

- 1-GB copper 12-port AppNav Controller Interface Module
- 1-GB SFP 12-port AppNav Controller Interface Module
- 10-GB SFP+ 4-port AppNav Controller Interface Module

AppNav Controller Interface Module interfaces are configured differently to support either in-path or off-path models of deployment:

- In-path—The ANC operates in inline interception mode with at least one inline bridge group configured on the AppNav Controller Interface Module. A bridge group consists of two or more physical or logical (port channel) interfaces.
- Off-path—The ANC operates in WCCP interception mode with one physical or logical (standby or port channel) interface configured with an IP address.

Interfaces on the AppNav Controller Interface Module can have three functions:

- Interception—Used to receive traffic intercepted from the network and egress traffic to the network. The interception interface is implied based on the AppNav Controller placement and does not require explicit configuration for this function.
- Distribution—Used to distribute traffic to the WNs and receive egressed traffic from the WNs. The distribution interface is explicitly configured as the cluster interface for intra-cluster traffic and must be assigned an IP address.
- Management—A management interface can be optionally and exclusively designated for management traffic and isolated from the normal data path. We recommend that you use one of the appliance's built-in interfaces for management traffic and reserve the high performance interfaces on the AppNav Controller Interface Module for interception and distribution.

You should use separate interfaces for interception and distribution for best performance, but you can use the same interface for both functions.

AppNav Controller Interface Modules support port channel and standby logical interfaces. A port channel allows you to increase the bandwidth of a link by combining multiple physical interfaces into a single logical interface. A standby interface allows you to designate a backup interface in case of a failure.

Interfaces on the AppNav Controller Interface Module support the following:

- A maximum of seven port channels with up to eight physical interfaces combined into a single port channel group.
- A maximum of five bridge groups configured over the physical or logical interfaces.

Interfaces on the AppNav Controller Interface Module do not support the following:

- Fail-to-wire capability
- Bridge virtual interfaces (BVIs)

AppNav Policy

The AppNav policy is a flow distribution policy that allows you to control how ANCs distribute traffic to the available WNs.

The AppNav policy consists of class maps that classify traffic according to one or more match conditions and a policy that contains rules that specify distribution actions to WNGs for each of the classes.

This section includes the following topics:

Class Maps

AppNav class maps classify traffic according to one or more of the following match conditions:

- Peer device ID—Matches traffic from one peer WAAS device, which could be handling traffic from a single site or a group of sites.

For example, you can use this kind of matching to classify all traffic from a peer device that serves one branch office.

- 3-tuple of source IP, and/or destination IP, and/or destination port (matches traffic from a specific application).

  For example, you can use this kind of matching to classify all HTTP traffic that uses port 80.

- A mix of one peer device ID and the source IP, and/or destination IP, and/or destination port (matches application-specific traffic from one site).

  For example, you can use this kind of matching to classify all HTTP traffic that is from a peer device that serves the one branch office.

The class-default class map is a system-defined default class map that is defined to match any traffic. By default, it is placed in the last rule in each policy to handle any traffic that is not matched by other classes.

Policies

An AppNav Controller matches incoming flows to class maps and the policy rules in a policy associate class maps with actions, such as distributing a flow to a particular WNG for optimization. The order in which rules are listed in the policy is important. Starting at the top of the policy, the first rule that matches a flow determines to which WNG it is distributed.

A policy rule can specify four kinds of actions to take on a flow:

- Specify the primary WNG to which to distribute the flow (required).
- Specify a backup WNG for distribution if the primary WNG is unavailable or overloaded (optional).

  The primary WNG receives all traffic until all WNs within the group become overloaded (reach 95 percent of the maximum number of transport flow optimization [TFO] connections) or are otherwise unavailable, and then traffic is distributed to the backup WNG. If a WN in the first WNG becomes available, traffic is again distributed there. If all WNs in both WNGs become overloaded, traffic is passed through unoptimized.

- Monitor the load on the application accelerator that corresponds to the application traffic matched by the class (optional).

  If the monitored application accelerator on one WN in a WNG becomes overloaded (reaches 95 percent of its maximum number of connections), the WN is considered overloaded and traffic is directed to another WN in the group. If all WNs become overloaded, traffic is distributed to the backup WNG. This application accelerator monitoring feature is useful for ensuring optimization for critical applications and is recommended for the MAPI and SMB accelerators.

- Specify a nested policy to apply to the flow (optional).

  For more information, see the "Nested Policies" section on page 4-6.

Within a WNG, flows are distributed evenly among WNs. If a WN reaches its maximum capacity or becomes unavailable, it is not sent new flows. New flows are sent to other available WNs in the WNG so that they can be optimized successfully.

Note If a WN that is doing MAPI or ICA application acceleration becomes overloaded, flows associated with existing MAPI and ICA sessions continue to be sent to the same WN due to the requirement that the same WN handle these types of flows. New MAPI and ICA flows, however, are distributed to other WNs.

The AppNav policy is specific to each ANC, though typically all ANCs in a cluster have the same policy. Each ANC consults its AppNav policy to determine which WNG to use for a given flow. Different ANCs in a cluster can have different AppNav policies, which allows you to customize distribution in certain cases. For example, when a cluster contains ANCs and WNs that are in different locations, it may be more desirable for an ANC to distribute traffic to WNs that are closer to it.

Nested Policies

A policy rule can specify one nested policy, which allows traffic identified in a class to be subdivided and handled differently. Nested policies provide two advantages:

- It allows another policy to be used as a common subclassification tool.

  For example, you can define a policy that contains monitoring actions and apply it as a subpolicy to multiple classes in the primary policy.
- It provides a method of including class maps with both match-any and match-all characteristics into a single subclass.

The nested policy feature is designed for use with site-based classes (matched by peer ID) at the first level and application-based subclasses (matched by IP address/port) at the second level. Only the first level policy can contain classes that use match peer conditions.

Site and Application Affinity

You can provision a WNG for serving specific peer locations (site affinity) or applications (application affinity) or a combination of the two. Using a WNG for site or application affinity provides the following advantages:

- Provisioning—Localize a class of traffic to achieve control over provisioning and performance monitoring. For example, a business-critical application like Sharepoint or a business-critical site can be given assured capacity and monitored closely for performance.
- Enhanced application performance—Better compression performance is achieved by limiting data that belongs to a site to one or a few WNs, which results in better utilization of the Data Redundancy Elimination (DRE) cache.

Figure 4-3 depicts how sites and applications can be associated with node groups. The following WNGs are defined:

- WNG-1—Consists of two WNs that process flows coming only from sites A and B.
- WNG-2—Consists of two WNs that process HTTP and SSL flows from any site. Whether HTTP and SSL flows from Site A and Site B should be processed by WNG-2 or WNG-1 is determined by the order of rules in the policy.
- WNG-3—Consists of two WNs that process MAPI flows coming from any site. Whether MAPI flows from Site A and Site B should be processed by WNG-3 or WNG-1 is determined by the order of rules in the policy.
- WNG-4—Consists of three WNs. The class-default class is applied to this WNG so that it is sent all flows that do not match any other class map.

Figure 4-3 Flow Distribution Using Site and Application Affinity

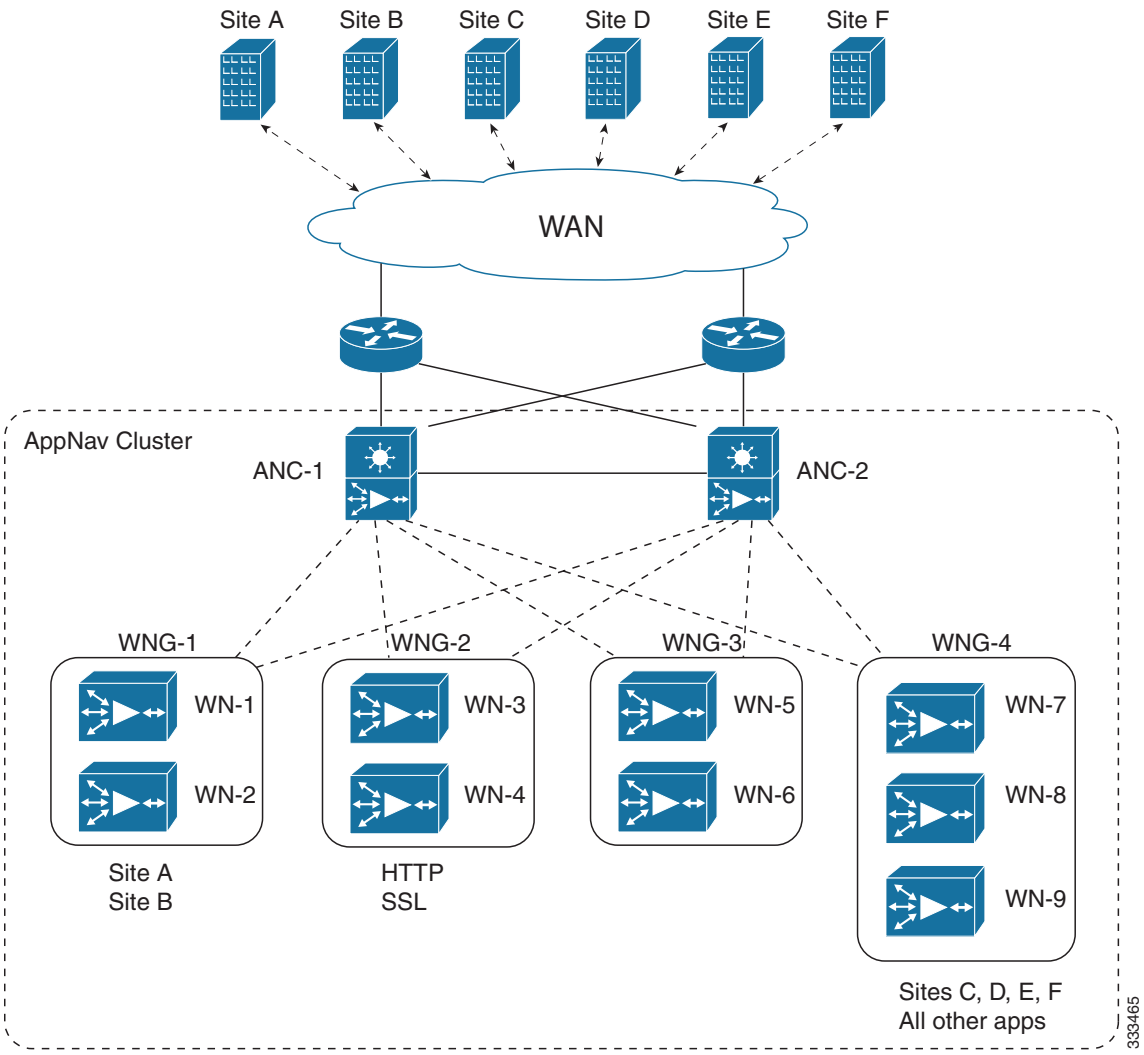

The following sections provide more details about these topics:

- Site Affinity, page 4-7
- Application Affinity, page 4-8

Site Affinity

Site affinity gives you the ability to always send all traffic from one site to a specific WNG, which allows you to reserve optimization capacity for critical sites and to improve compression performance through better utilization of the DRE cache.

Traffic from any location, not just a single site, can be matched in a class map and associated with a WNG.

You can implement site affinity by configuring a class map that matches the device ID of the WAE in the site. If a site has more than one WAE in a WCCP farm or a serial inline cluster, specify multiple device IDs in the class map. Next, associate the class map with a distribution action to a WNG in a policy rule.

You can also identify sites using source IP addresses or subnets in the class map, if you know what IP addresses are used in the site and keep the policy configuration consistent with site IP addresses. However, we recommend that you use peer device IDs in configuring site affinity.

Note A peer ID-based class map works only for matching flows that carry the WAAS auto discovery TCP options. If you configure a class to match a site peer ID at the data center, the same class does not match flows that originate in the other direction, such as those flows that originate from the data center and go back to the same site. Such flows are usually small in number compared to the site to data center flows.

If you want flows in both directions to go to the same WNG, you must configure two class maps: one to match in the site to data center direction, typically using the site device ID; and another to match the data center to site direction, using destination IP subnets belonging to the site. Both class maps can be configured to distribute traffic to the same WNG. A mesh network is a specific use case where flows can originate in either direction.

If the site WAE is in overload or does not mark the SYN packet with auto discovery options for any other reason, the ANC cannot match it to the peer match class map.

Application Affinity

Application affinity gives you the ability to always send certain application traffic to a specific WNG, which allows you to reserve optimization capacity for different applications depending on business priorities.

In the context of AppNav flow distribution, an application is defined using a three-tuple of the source IP, destination IP, and destination TCP port. The actual type of traffic does not matter for flow distribution. For example, you can use separate WNGs for HTTP traffic that is addressed to different destination ports or different server IP addresses. Destination IP and ports are most useful in using application affinity, but having the source IP also helps you to define the traffic of interest.

A small number of protocols, such as FTP, use dynamic destination ports. An FTP server in active mode originates a data connection back to the FTP client using a dynamic destination port. This port is exchanged over the control channel from client to server using the well-defined destination port 21. Consider trying to define a class map for FTP. Because the destination port is not known in advance, you cannot map both control and data connections to the same class. In this case, we recommend that you use the client IP addresses or subnets to match against destination IP addresses for the data connections. You must configure two class maps: one for the control channel, using destination port 21, and another for the data channel, using destination IP addresses. You can configure policy rules so that both class maps distribute traffic to the same WNG.

You can further classify traffic from a site into applications by combining the peer matches with three-tuple matches in a match-all class map, called a Custom class map type in the Central Manager. You can define separate WNGs, for example, for HTTP traffic from a particular site and CIFS traffic from the same site.

Default Policy Behavior

The following default class maps are provided:

- CIFS—Matches traffic for destination ports 139 and 445
- Citrix-ICA—Matches traffic for destination port 1494
- Citrix-CGP—Matches traffic for destination port 2598

- epmap—Matches traffic for destination port 135
- HTTP—Matches traffic for destination ports 80, 3128, 8000, 8080, and 8088
- HTTPS—Matches traffic for destination port 443
- MAPI—Matches traffic for the MS RPC MAPI application (dynamic port assignment)
- NFS—Matches traffic for destination port 2049
- RTSP—Matches traffic for destination ports 554 and 8554
- class-default—Matches any TCP traffic (this class map cannot be edited or deleted)

If you use the Central Manager AppNav Cluster Wizard to create an AppNav Cluster, the wizard creates a default policy named appnav_default. This policy is assigned by default to all ANCs in a cluster and contains only the class-default policy rule that has the following characteristics:

- Matches class-default (any TCP) traffic.
- Distributes class-default traffic to the default WNG, which includes all WNs created by the wizard, with no backup WNG specified.
- Contains the waas_app_default nested policy, which provides application monitoring for each of the default class maps, except video (RTSP).

  When you use the Central Manager to define a policy rule for any class that uses peer matching or source or destination IP address matching (but not port matching), it automatically adds the waas_app_default policy as a nested policy. The waas_app_default policy is created by the system and monitors all application accelerators (except video), so you do not need to manually add application accelerator monitoring to your policy rules, unless it is for the video accelerator.

If you do not use the Central Manager AppNav Cluster Wizard to create a cluster, there is no default flow distribution, so if an incoming flow does not match any class in the AppNav policy, it is not distributed to any WNG; instead, it is passed through.

If a WNG is defined but is not used in any policy rule, it does not receive any flows. If a policy is defined but not applied to an ANC, it does not take effect.

The default action for a policy rule is none, which is context dependent: in a top level policy it means pass through and if the policy is nested, it means inherit the parent policy rule action.

Prerequisites for AppNav Deployment

AppNav has the following prerequisites:

- Each WAAS appliance to be used as an AppNav Controller must contain a Cisco AppNav Controller Interface Module.
- Each AppNav Controller must be configured in appnav-controller device mode.

Guidelines and Limitations

AppNav has the following configuration guidelines and limitations:

- An AppNav Cluster can contain a maximum of the following:
  - 8 ANCs
  - 32 WNs

  - 32 WNGs
- All ANCs in an ANCG must have the same set of ANCs and WNGs in their configuration.
- All WNs in one WNG must have identical optimization policies configured on them.
- AppNav class maps and policies can be configured only at the cluster level, not at the device level, from the Central Manager. At the device level, class maps and policies may only be viewed.
- You can define the following maximum policy entities within a service context:
  - 1024 match conditions
  - 512 AppNav class maps
  - 64 rules per AppNav policy
  - 64 AppNav policies, though only one policy is actively bound to the service context and used for flow distribution on a given ANC
- There is no fail-to-wire capability on AppNav Controller Interface Module interfaces configured in bridge groups for inline mode, which would allow traffic to bypass the interface if the device fails or loses power. Therefore, if you are using inline mode, we recommend that you deploy two or more AppNav Controller appliances to provide high availability.
- Virtual blades are not supported on WAAS appliances that are operating as AppNav Controllers.

Configuring an AppNav Cluster

This section contains the following topics:

Task Flow for Configuring an AppNav Cluster

You must complete the following steps to configure an AppNav Cluster:

1. Install and configure the individual ANC and WN devices with basic network settings. See the "Configuring WAAS Device Interfaces" section on page 4-11.
2. Use the Central Manager AppNav Cluster Wizard to create a cluster and configure the interception mode, configure cluster settings, choose cluster devices, configure traffic interfaces, and configure WCCP settings if you are using WCCP. See the "Creating a New AppNav Cluster with the Wizard" section on page 4-14.

3. (Optional) Configure AppNav class maps. This step is necessary only if you want to customize the default class map configuration. The system adds several default class maps that match traffic corresponding to most of the application accelerators and a class-default class map that matches all traffic. See the "Configuring AppNav Class Maps" section on page 4-19.
4. (Optional) Configure an AppNav policy. This step is necessary only if you want to customize the default policy. The system adds a default policy that distributes all traffic to the WNG-Default WNG, which is the node group into which all WNs are grouped by default. See the "Configuring Rules Within an AppNav Policy" section on page 4-22.
5. (Optional) Configure WAAS node optimization class maps and policy rules. This step is necessary only if you want to customize the default optimization policy that is listed in Appendix A, "Predefined Optimization Policy."
6. (Optional) Configure an interception ACL on the ANCs. See the "Configuring AppNav Controller ACLs" section on page 4-26.

Configuring WAAS Device Interfaces

Before you can use the AppNav Cluster wizard to create an AppNav Cluster, you must connect the WAAS device interfaces and configure the management interfaces. Configuration differs depending on whether management traffic uses a separate interface or shares the traffic handling interface.

This section contains the following topics:

For more information about device interface configuration, see Chapter 6, "Configuring Network Settings." For more information about configuring a bridge group for inline interception mode, see the "Configuring Inline Operation on ANCs" section on page 5-49.

Interface Configuration with a Separate Management Interface

If you want management traffic to use a dedicated interface, separate from the traffic data path, connect and configure the devices as described in this section.

AppNav Controller

Step 1 Connect the last AppNav Controller Interface Module port to the switch/router port for the cluster traffic. For example, this port is GigabitEthernet 1/11 on a 12-port module or TenGigabitEthernet 1/3 on a 4-port module.

Step 2 Connect a built-in Ethernet port to the switch/router port for the management interface.

Step 3 For an in-path (inline) deployment, connect the first pair of ports on the AppNav Controller Interface Module (for example, GigabitEthernet 1/0 [LAN] and GigabitEthernet 1/1 [WAN] for bridge 1) to corresponding switch/router ports.

If the ANC is connected to a second router for a dual inline deployment, connect the second pair of ports on the AppNav Controller Interface Module (for example, GigabitEthernet 1/2 [LAN] and GigabitEthernet 1/3 [WAN] for bridge 2) to corresponding switch/router ports.

Step 4 Use the device **setup** command to configure the following settings:

- Configure the device mode as AppNav Controller.
- Configure the IP address and netmask of the built-in management port.
- Configure the built-in management port as the primary interface.
- Configure the other network and basic settings (default gateway, DNS, NTP server, and so forth).
- Register the device with the Central Manager by entering the Central Manager IP address.

Step 5 Configure the IP address and netmask of the last AppNav Controller Interface Module port. You can also configure these settings through the AppNav Cluster wizard, if desired.

WAAS Node

Step 1 Connect a built-in Ethernet port to the switch/router port for the management interface.

Step 2 Use the device **setup** command to configure the following settings:

- Configure the device mode as Application Accelerator.
- Configure the IP address and netmask of the built-in management port.
- Configure the built-in management port as the primary interface.
- Configure the other network and basic settings (default gateway, DNS, NTP server, and so forth).
- Register the device with the Central Manager by entering the Central Manager IP address.

Interface Configuration with a Shared Management Interface

If you want management traffic to use an interface shared by the traffic data path, connect and configure the devices as described in this section.

AppNav Controller

Step 1 Connect the last AppNav Controller Interface Module port to the switch/router port for the cluster traffic. For example, this port is GigabitEthernet 1/11 on a 12-port module or TenGigabitEthernet 1/3 on a 4-port module.

Step 2 For an in-path (inline) deployment, connect the first pair of ports on the AppNav Controller Interface Module (for example, GigabitEthernet 1/0 [LAN] and GigabitEthernet 1/1 [WAN] for bridge 1) to corresponding switch/router ports.

If the ANC is connected to a second router for a dual inline deployment, connect the second pair of ports on the AppNav Controller Interface Module (for example, GigabitEthernet 1/2 [LAN] and GigabitEthernet 1/3 [WAN] for bridge 2) to corresponding switch/router ports.

Step 3 Use the device **setup** command to configure the following settings:

- Configure the device mode as AppNav Controller.
- Configure the IP address and netmask of the last AppNav Controller Interface Module port.
- Configure the last AppNav Controller Interface Module port as the primary interface.

- Configure the other network and basic settings (default gateway, DNS, NTP server, and so forth).
- Register the device with the Central Manager by entering the Central Manager IP address.

WAAS Node

Step 1 Connect a built-in Ethernet port to the switch/router port for the management interface.

Step 2 Use the device **setup** command to configure the following settings:

- Configure the device mode as Application Accelerator.
- Configure the IP address and netmask of the built-in management port.
- Configure the built-in management port as the primary interface.
- Configure the other network and basic settings (default gateway, DNS, NTP server, and so forth).
- Register the device with the Central Manager by entering the Central Manager IP address.

Interface Configuration Considerations

The following guidelines concern WAAS device interface configuration:

- On an ANC, the intercepted traffic must go through an interface on the AppNav Controller Interface Module.
- On an ANC that also serves as a WN, the cluster interface is the same as the interception interface.
- On a WN, cluster traffic can be handled on any interface, either built-in or on an interface module.
- To simplify AppNav deployment, the AppNav Cluster Wizard uses the following conventions for configuring the AppNav Controller Interface Module ports on an ANC:
  - The default port for cluster traffic is the last port on the module (for example, GigabitEthernet 1/11 on a 12-port module or TenGigabitEthernet 1/3 on a 4-port module).
  - For an in-path (inline) deployment, the default interception bridge is the first pair of ports on the module (for example, GigabitEthernet 1/0 [LAN] and GigabitEthernet 1/1 [WAN] for bridge 1). If the ANC is connected to a second router for a dual inline deployment, the default second interception bridge is the second pair of ports on the module (for example, GigabitEthernet 1/2 [LAN] and GigabitEthernet 1/3 [WAN] for bridge 2).

The AppNav Cluster Wizard uses four predefined deployment models to help simplify configuration. Each deployment model expects interfaces to be connected and configured in a particular way, except for the Custom option, which allows you to configure interfaces in any way. Before you run the wizard with one of the four predefined models, the needed interfaces must be in either of these states:

- Not configured with an IP address and netmask and not used as part of another logical interface. (However, the last port on the AppNav Controller Interface Module can be configured with an IP address because it is the default port for cluster traffic.)

  The wizard configures all needed traffic interface settings.
- Configured as expected by the wizard according to the following deployment model expectations.

The following sections describe the interface configurations used by each of the four predefined deployment models.

Single AppNav Controller WCCP Interception

With a 12-port AppNav Controller Interface Module:

- Port channel 1—Contains ports GigabitEthernet 1/10 and 1/11
- Cluster interface—Port channel 1

With a 4-port AppNav Controller Interface Module:

- Cluster interface—GigabitEthernet 1/3

Dual AppNav Controllers WCCP Interception

With a 12-port AppNav Controller Interface Module:

- Port channel 1—Contains ports GigabitEthernet 1/10 and 1/11
- Port channel 2—Contains ports GigabitEthernet 1/8 and 1/9
- Standby group 1—Contains interfaces Port channel 1 (primary) and Port channel 2
- Cluster interface—Standby Group 1

With a 4-port AppNav Controller Interface Module:

- Standby group 1—Contains ports GigabitEthernet 1/2 and 1/3 (primary)
- Cluster interface—Standby Group 1

Single AppNav Controller Inline Interception

- Interception bridge 1—Contains ports GigabitEthernet 1/0 (LAN) and 1/1 (WAN)
- Cluster interface—GigabitEthernet 1/11

Dual AppNav Controllers Inline Interception

- Interception bridge 1—Contains ports GigabitEthernet 1/0 (LAN) and 1/1 (WAN)
- Interception bridge 2—Contains ports GigabitEthernet 1/2 (LAN) and 1/3 (WAN)
- Standby group 1—Contains ports GigabitEthernet 1/10 and 1/11 (primary)
- Cluster interface—Standby Group 1

Creating a New AppNav Cluster with the Wizard

Prerequisites

- Set up the individual ANC and WN devices as described in the "Configuring WAAS Device Interfaces" section on page 4-11.
- Ensure that all ANCs are configured for AppNav Controller device mode. If you need to change the device mode, see the "Changing Device Mode" section on page 2-16.
- Use the Central Manager to configure basic settings for all devices such as NTP server, AAA, logging, and so on.

Detailed Steps

To create a new AppNav Cluster by using the wizard, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **AppNav Clusters > All AppNav Clusters**.

The Manage AppNav Clusters window appears.

Step 2 Click the **AppNav Cluster Wizard** icon in the taskbar of the Manage AppNav Clusters area. The Cluster Wizard window appears.

Step 3 In the Deployment model drop-down list, choose one of the following deployment models that matches your deployment:

- **Single AppNav Controller WCCP interception**
- **Dual AppNav Controllers WCCP interception**
- **Single AppNav Controller Inline interception**
- **Dual AppNav Controllers Inline interception**
- **Custom**—For a deployment that does not match one of the choices above

Click **Next**.

Step 4 (Optional) If you chose the Custom deployment model, from the Interception method drop-down list, choose the **WCCP** or **Inline interception** method and click **Next**.

Step 5 Define the cluster settings by entering the following information:

- In the Name field, enter a name for the cluster. Use only letters, numbers, hyphen, and underscore, up to a maximum of 32 characters and beginning with a letter.
- (Optional) In the Description field, enter a description of the cluster. Use only letters and numbers, up to a maximum of 200 characters.
- Check the **Disable Distribution** check box if you want make the cluster operate in monitoring mode, otherwise, it is activated when the wizard finishes. In monitoring mode, all traffic is passed through instead of being distributed to WNs.

Click **Next**.

Step 6 Choose the ANC and WN devices that you want to be part of the cluster:

a. Choose up to eight ANCs in the AppNav Controller device list by clicking the check box next to the device names. You can use the filter settings in the taskbar to filter the device list.

b. (Optional) If you want to enable optimization on the ANC devices, check the **Enable WAN optimization on selected AppNav Controller(s)** check box (it may be enabled or disabled by default, depending on the deployment model you chose).

c. Choose up to 32 WNs in the WAAS Nodes device list by clicking the check box next to the device names. You can use the filter settings in the taskbar to filter the device list.

If there are devices that are ineligible to join the cluster, click **Show Ineligible Devices** to see them and the reasons why they are ineligible. You can use the filter settings to filter the list.

d. Click **Next**.

Step 7 Verify the cluster interface, IP address, and netmask for each device in the cluster. The wizard automatically selects recommended cluster interfaces that should be configured. To edit the IP address and netmask settings for a device, choose the device and click the **Edit** taskbar icon. This screen does not appear if you are configuring a custom cluster.

Click **Finish** if you are using inline interception (and you are done) or click **Next** if you are using WCCP interception (and continue with the following steps for WCCP).

Step 8 (Optional) Configure the WCCP settings for the ANC. This screen does not appear if you are configuring an inline cluster.

For details about configuring WCCP, see the "Configuring WCCP on WAEs" section on page 5-11.

a. Ensure the **Enable WCCP Service** check box is checked if you want to enable WCCP. This item appears only if you are defining a custom cluster.

b. Verify the single WCCP service ID of 61 (default) or change it if desired.

   You need to configure only this single WCCP service on both the ingress and egress ports of the router doing WCCP redirection to this ANC.

c. (Optional) If you want to enable two WCCP services, uncheck the **Enable Single Service Mode** check box (it is checked by default because two WCCP services are not needed). The automatically assigned second service ID number is shown in the Service ID2 field.

d. From the Redirect Method drop-down list, choose the WCCP L2 or WCCP GRE redirect method. For details on the redirect method, see the "Configuring or Viewing the WCCP Settings on ANCs" section on page 5-22. This item appears only if you are defining a custom cluster.

e. (Optional) If you do not want to use the default gateway defined on the device, uncheck the **Use Default Gateway as WCCP Router** check box. Enter the address of one or more WCCP routers, separated by commas, in the WCCP Routers field.

f. Click **Advanced WCCP Settings** to configure additional settings as needed. For more information on these fields, see the "Configuring or Viewing the WCCP Settings on ANCs" section on page 5-22. This item appears only if you are defining a custom cluster.

g. Click **Next**. If you are configuring multiple ANCs, a similar screen is shown for each ANC.

Step 9 Configure the interception and cluster interface settings for each device. The Cluster Interface Wizard appears only if you are defining a custom cluster, with one screen for each device in the cluster:

a. Configure individual interfaces, port channels, standby interfaces, and bridge interfaces (for inline only) as needed on the device by using the graphical interface wizard. If you are configuring an inline ANC, you must define a bridge interface with two physical or port-channel interfaces (or one of each). For details on how to use the wizard, see the "Configuring Interfaces with the Graphical Interface Wizard" section on page 4-17.

b. From the Cluster Interface drop-down list, choose the interface to be used for intra-cluster traffic.

c. Click **Next**. If you are configuring multiple devices, a similar screen is shown for each device.

Step 10 Click **Finish** to save the cluster configuration.

By default, the wizard assigns all WNs to a default WNG named WNG-Default. You can create additional WNGs as described in the "Adding a New WAAS Node to the Cluster" section on page 4-32. You can reassign WNs to different WNGs as described in the "Configuring WAAS Node Settings" section on page 4-29.

After you create an AppNav Cluster, it is shown in the Manage AppNav Clusters list. For details on monitoring the cluster, see the "Monitoring an AppNav Cluster" section on page 4-34.

Configuring Interfaces with the Graphical Interface Wizard

You can easily configure interfaces on AppNav Controller Interface Modules that are installed in devices that are part of an AppNav Cluster by using the graphical interface wizard (see Figure 4-4).

Figure 4-4 Graphical Interface Wizard

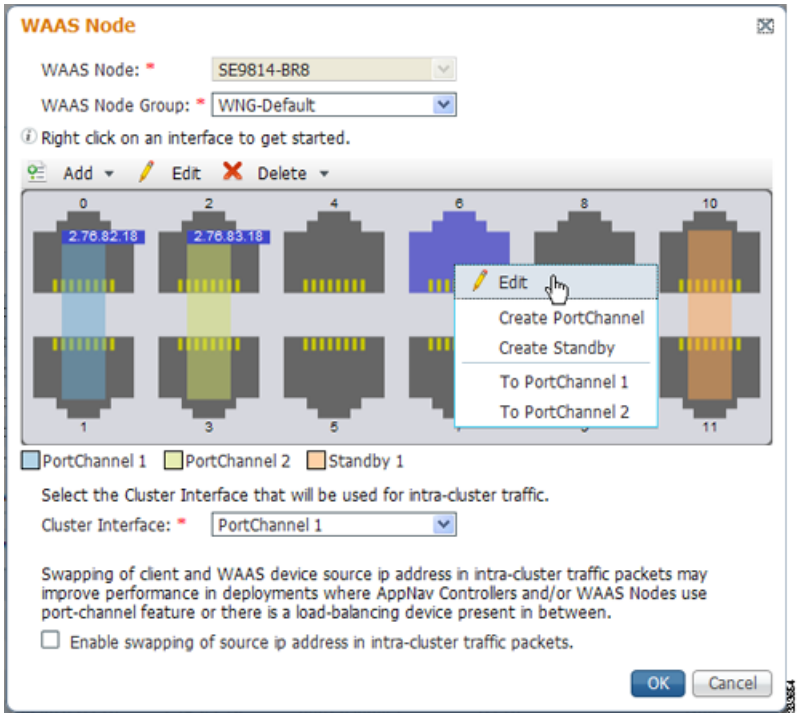

The graphical interface wizard appears when you are editing the settings for a WN or ANC in the AppNav Cluster context. The top two fields, WAAS Node and WAAS Node Group, do not appear when configuring ANC interfaces.

In the graphical interface view, hover over a physical or logical interface to see its identifier (for example, GigabitEthernet 1/0). Port channels, bridge groups, and standby groups are indicated by colored blocks or dotted outlines. The IP address of each configured physical or logical interface is shown in a small blue highlight. The legend below the table indicates port channel, bridge group, and standby interfaces.

Right click on an interface to choose from the following actions:

- Edit—To display a pane where you can edit the interface description, IP address, netmask, and shutdown status.
- Create PortChannel—To create a new port channel with this interface. This choice displays a pane where you can configure the port channel number, description, IP address, netmask, and shutdown status.
- Create Bridge—To create a new bridge group with this interface. This choice displays a pane where you can configure the bridge group number and description and enable link state propagation. This choice appears only when configuring a device for inline interception. A bridge interface consists of two physical or port-channel interfaces (or one of each)

- Create Standby—To create a new standby group with this interface. This choice displays a pane where you can configure the standby group number, description, IP address, netmask, and shutdown status.
- To PortChannel *n*—To add this interface to an existing port channel, where *n* is the port channel number.
- To Standby *n*—To add this interface to an existing standby group, where *n* is the standby group number.
- To Bridge *n*—To add this interface to an existing bridge group, where *n* is the bridge group number.
- For standby interfaces (right-click within the standby interface group indicator):
  - Edit—To edit the standby group settings such as the description, IP address, netmask, primary interface, and shutdown status.
  - Delete Standby *n*—To delete the standby group.
- For port channel interfaces (right-click within the port channel indicator):
  - Edit—To edit the port channel settings such as the port channel number, description, IP address, netmask, and shutdown status.
  - Remove from Standby *n*—To remove the port channel from standby group *n*.
  - Delete PortChannel *n*—To delete the port channel.
- For bridge group interfaces (right-click within the bridge group indicator):
  - Edit—To edit the bridge group settings such as the bridge group number, description, and link state propagation status.
  - Delete Bridge *n*—To delete the standby group.

To select an interface:

- Individual interface—Click and selection is indicated by a blue color.
- Standby group—Click on colored or dotted line indicator and selection is indicated by a thick dotted blue outline around all interfaces in the standby group.
- Port channel or bridge group—Click on colored indicator and selection is indicated by a thick dotted blue outline around all interfaces in the port channel or bridge group.

You can also perform actions by selecting an interface and clicking the following taskbar icons:

- Add (choices differ depending on the selected entity):
  - Create PortChannel—To create a new port channel with this interface.
  - Create Bridge—To create a new bridge group with this interface.
  - Create Standby—To create a new standby group with this interface.
  - To PortChannel *n*—To add this interface to an existing port channel, where *n* is the port channel number.
  - To Standby *n*—To add this interface to an existing port channel, where *n* is the port channel number.
- Edit—To edit the selected interface.
- Delete (choices differ depending on the selected entity):
  - Remove from Standby *n*—To remove the port channel from standby group *n*.
  - Delete PortChannel *n*—To delete the port channel.
  - Delete Standby *n*—To delete the standby group.

– Delete Bridge *n*—To delete the bridge group.

Use the Cluster Interface drop-down list to select the interface to be used for intra-cluster traffic (between the ANCs and WNs).

To enable swapping of client and WAAS device source IP address fields in intra-cluster traffic, check the **Enable swapping of source IP address in intra-cluster traffic** check box. You may want to enable this option if you are using a port channel for the cluster interface or there is a load balancing device between the ANC and WN. This option may improve the load balancing of traffic that the ANC distributes to WNs for optimization because it load balances based on the client IP address rather than the ANC IP address. (For traffic from the server to the client, it swaps the server IP address with the ANC IP address.)

Note If you are using WCCP, the WCCP control messages must pass through the ANC interface that receives intercepted traffic from the routers. If WCCP control messages are routed to the ANC management interface, the cluster does not operate.

Configuring AppNav Policies

This section contains the following topics:

Configuring AppNav Class Maps

To configure AppNav class maps, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **AppNav Clusters >** *cluster-name*.

Step 2 Choose **Configure > AppNav Cluster > AppNav Class-Map**.

The AppNav Class-Maps window appears, listing the existing class maps.

From this window, you can perform the following tasks:

- Use the filter settings in the Show drop-down list to filter the class map list as needed. You can use a quick filter or show all class maps.
- Edit a class map by selecting it and clicking the **Edit** taskbar icon.
- Delete one or more class maps by selecting them and clicking the **Delete** taskbar icon.
- Add a new class map as described in the steps that follow.

Step 3 Click the **Add Class-Map** taskbar icon.

Step 4 In the Name field enter a name for the class map.

Step 5 (Optional) In the Description field enter a description for the class map.

Step 6 From the Type drop-down list, choose the class map type:

- **Application**—Matches traffic for a particular application based on source and/or destination IP addresses and/or ports, or the Microsoft RPC application identifier (for applications that use dynamic port allocation). Continue with Step 7.

- **Site**—Matches traffic from particular WAAS peer devices, for site affinity. Continue with Step 8.
- **Custom**—Mixes application and site affinity. Matches traffic for a particular application from one specific peer WAAS device. Continue with Step 9.
- **Any TCP**—Matches any TCP traffic as a catch-all classifier. If you choose this type, there are no other fields to set. Click **OK** to finish and return to the class maps list.

The match conditions shown in the lower part of the pane change depending on the class map type.

Step 7 (Optional) For an Application class map type, enter one or more match conditions. You can perform the following tasks in this pane:

- Edit a match condition by selecting it and clicking the **Edit** taskbar icon.
- Delete one or more match conditions by selecting them and clicking the **Delete** taskbar icon.
- Add a new match condition as described in the steps that follow.

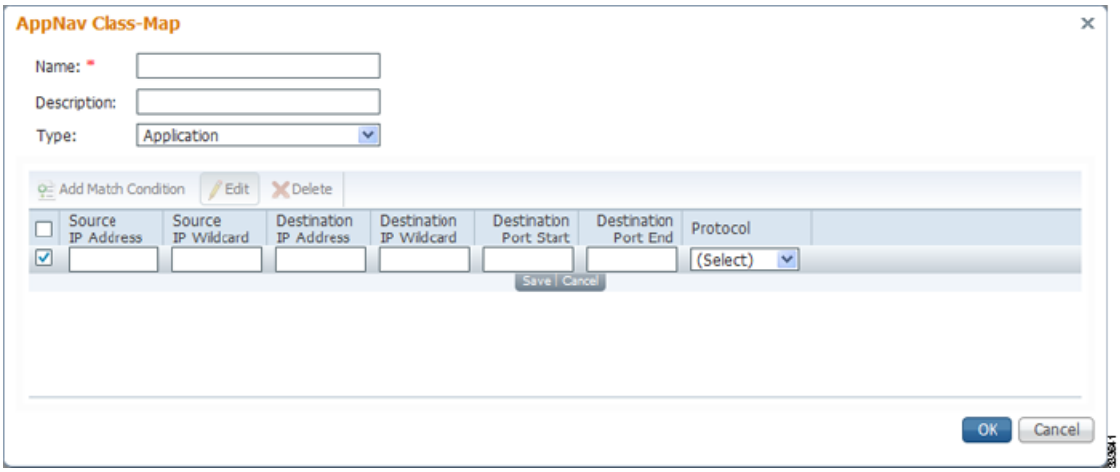

a. Click the **Add Match Condition** taskbar icon.

b. Enter values in one or more fields to create a condition for a specific type of traffic. For example, to match all traffic going to ports 5405–5407, enter **5405** in the Destination Port Start field and **5407** in the Destination Port End field. You can use the IP address wildcard fields to specify a range of IP addresses using a wildcard subnet mask in dotted decimal notation (such as 0.0.0.255 for /24).

c. If you want to match Microsoft RPC traffic that uses dynamic port allocation, choose the RPC application identifier from the Protocol drop-down list. For example, to match Microsoft Exchange Server traffic that uses the MAPI protocol, choose **mapi**.

d. Click **Save** to save the match condition.

e. Add additional match conditions as needed and click **OK** to save the class map and return to the class maps list. If any one of the conditions is matched, the class is considered matched.

Step 8 (Optional) For a Site class map type, select one or more peer devices. Follow these steps to create the class map:

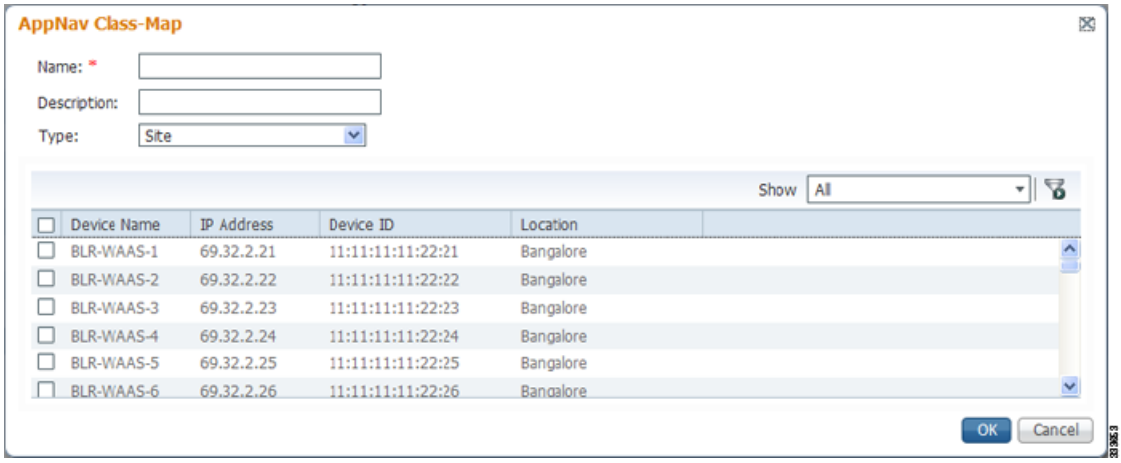

a. Use the filter settings in the Show drop-down list to filter the device list as needed. You can use a quick filter, show all devices, or show all assigned devices.

b. Check the box next to each device that you want to match traffic from. You can check the box next to the column titles to select all devices and uncheck it to deselect all devices. If any one of the selected devices is matched, the class is considered matched.

c. Click **OK** to save the class map and return to the class maps list.

Step 9 (Optional) For a Custom class map type, you must enter one match condition based on IP address/port or Microsoft RPC application ID and you must choose one WAAS peer device. All specified matching criteria must be satisfied for the class to be considered matched. Follow these steps to create the class map:

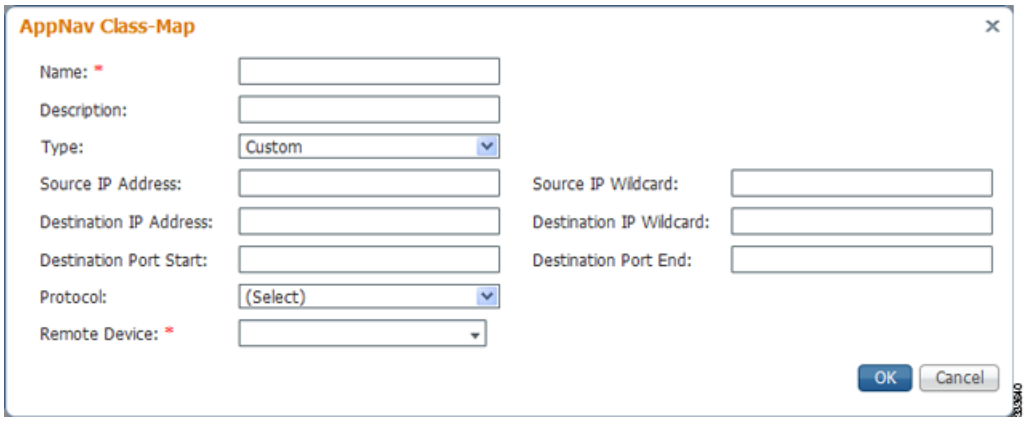

a. Enter values in one or more IP address and/or port fields to create a condition for a specific type of traffic. For example, to match all traffic going to ports 5405–5407, enter **5405** in the Destination Port Start field and **5407** in the Destination Port End field. You can use the IP address wildcard fields to specify a range of IP addresses using a wildcard subnet mask in dotted decimal notation (such as 0.0.0.255 for /24).

b. (Optional) If you want to match Microsoft RPC traffic that uses dynamic port allocation, choose the RPC application identifier from the Protocol drop-down list. For example, to match Microsoft Exchange Server traffic that uses the MAPI protocol, choose **mapi**.

c. You must choose one WAAS peer device from the Remote Device drop-down list.

d. Click **OK** to save the class map and return to the class maps configuration window.

Configuring Rules Within an AppNav Policy

To configure rules in an AppNav policy, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **AppNav Clusters >** *cluster-name*.

Step 2 Choose **Configure > AppNav Cluster > AppNav Policies**.

The AppNav Policy window appears.

Step 3 Choose the policy to configure from the **AppNav Policy** drop-down list at the top.

You can click **Manage** to create or delete a policy or configure the ANCs to which a policy is applied. For details see the "Managing AppNav Policies" section on page 4-24.

From the AppNav Policy Rules area, you can perform the following tasks:

- Use the filter settings in the Show drop-down list to filter the rule list as needed. You can use a quick filter or show all rules.
- Edit a rule by selecting it and clicking the **Edit** taskbar icon.
- Delete one or more rules by selecting them and clicking the **Delete** taskbar icon.
- Move one or more selected rules to a new position by clicking the **Move To** taskbar icon. After moving the rows, click **Save Moved Rows** to save the change.
- Move one or more selected rules up or down one position by clicking the **Up** or **Down Arrow** taskbar icons, then click **Save Moved Rows** to save the change.
- Save rows that you have moved with the Move To or Up and Down Arrow functions by clicking the **Save Moved Rows** taskbar icon.
- Insert a new rule before the selected row by clicking the **Insert** taskbar icon. The workflow for inserting is the same as for adding (described in the following steps).
- Add a new rule at the end of the list as described in the steps that follow. (The class-default rule is always pushed to the last position.)

Step 4 Click the **Add Policy Rule** taskbar icon.

Step 5 From the AppNav Class-Map drop-down list, choose the class map to which this policy rule applies.

If you want to edit the class map, click **Edit**, or if you want to create a new class map, click **Create New**. The workflow is the same as described in the "Configuring AppNav Class Maps" section on page 4-19.

Step 6 From the Distribute To drop-down list, choose the distribution action to apply to the class map. The list includes all defined WNGs and the choices (None), for no action, and (Passthrough), to pass through this type of traffic. The meaning of (None) is context dependent: in a top level policy it means pass through and if this policy is nested, it means inherit the parent policy rule action.

When you choose a WNG, other settings appear. If you want create a new WNG, click **Create New**. The workflow is the same as described in the "Adding a New WAAS Node Group to the Cluster" section on page 4-34. The newly created WNG appears in both the Distribute To and Backup drop-down lists.

Step 7 (Optional) From the Backup drop-down list, choose the backup WNG to use for distribution if the primary WNG is unavailable.

Step 8 (Optional) From the Monitor drop-down list, choose the application accelerator to monitor. When you monitor an application accelerator, the ANC checks for overload on that application accelerator and does not send new flows to a WN that is overloaded. If you choose None, a specific application accelerator is not monitored, only the maximum connection limit of the device is monitored.

Step 9 (Optional) If you want to apply a nested policy within this rule, click **Nested Actions (Advanced)** to expand this area.

Step 10 (Optional) From the Nested Policy drop-down list, choose the policy to nest, or choose **None** to select no policy. When you choose a policy, the policy rules are displayed in a table.

If there are policies that are ineligible to be specified as a nested policy, click **Show Ineligible Policies** to display them and the reasons they are ineligible. A policy is ineligible if it already has a nested policy, because only one level of nesting is allowed.

To edit the chosen policy, click **Edit**, or to create a new policy for nesting, click **Create New**. The workflow for both editing and creating is the same.

a. In the Name field enter the policy name. This field is not editable for the waas_app_default policy.

b. Click the **Add Policy Rule** taskbar icon.

   A new row is added, showing fields for configuring the rule.

c. From the Class-Map drop-down list, choose the class map to which this rule applies.

d. From the Distribute To drop-down list, choose the distribution action to apply to the class map. The list includes all defined WNGs and the choices (Inherit), to inherit this action from the parent policy, and (Passthrough), to pass through this type of traffic.

e. (Optional) From the Backup drop-down list, choose the backup WNG to use for distribution if the primary WNG is unavailable.

f. (Optional) From the Monitor drop-down list, choose the application accelerator to monitor.

g. Click **OK** to save the policy rule and return to the AppNav Policy Rule pane for the primary policy rule you are creating.

Step 11 Click **OK** to create the policy rule and return to the policy configuration window.

Note If all AppNav policies have been deleted and you add a new policy rule, the policy rule is added to a new appnav_default policy, which is created automatically.

Managing AppNav Policies

To create or delete AppNav policies or configure the ANCs to which policies apply, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **AppNav Clusters >** *cluster-name*.

Step 2 Choose **Configure > AppNav Cluster > AppNav Policies**.

The AppNav Policy window appears.

Step 3 Choose the policy to view from the AppNav Policy drop-down list at the top.

For details on using the AppNav Policy Rules area see the "Configuring Rules Within an AppNav Policy" section on page 4-22.

Step 4 Click **Manage**.

From the Manage AppNav Policies pane, you can perform the following tasks:

- Use the filter settings in the Show drop-down list to filter the policy list as needed. You can use a quick filter or show all policies.
- Edit a policy and configure the ANCs to which it applies by selecting it and clicking the **Edit** taskbar icon.
- Delete a policy by selecting it and clicking the **Delete** taskbar icon.
- Add a new policy as described in the steps that follow.

Step 5 Click the **Add Policy** taskbar icon.

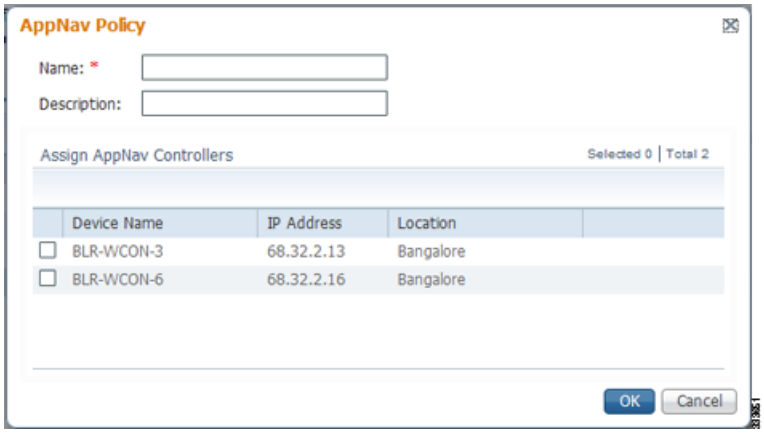

Step 6 In the Name field enter a name for the policy.

Step 7 (Optional) In the Description field enter a description for the policy.

Step 8 (Optional) Check the box next to each ANC that you want to assign to this policy. To unassign any assigned devices, uncheck the box.

Assigning the policy to an ANC makes the policy active on that ANC (only one policy can be active on an ANC) and removes the association of any previously active policy on that ANC. It is not necessary to assign the policy to an ANC if you want to create the policy as an alternate. You can assign it to ANCs later as needed.

Step 9 Click **OK** to save the policy and return to the Manage AppNav Policies pane.

Step 10 Click **Close** to return to the policy configuration window.

Step 11 Add policy rules to the new policy as described in the "Configuring Rules Within an AppNav Policy" section on page 4-22.

Configuring WAAS Node Optimization Policy

The WAAS node optimization policy controls how traffic that is distributed to the WAAS nodes is optimized. The optimization policy is configured on the WNs and any ANCs that are also acting as optimizing nodes.

All WNs in one WNG must have an identical optimization policy configured on them. Otherwise, optimization of flows is not predictable. The optimization policy can be different for different WNGs.

For information on how to configure the optimization policy, see Chapter 13, "Configuring Application Acceleration."

The default optimization policy is listed in Appendix A, "Predefined Optimization Policy."

Configuring AppNav Controller ACLs

An AppNav Controller ACL controls what traffic is intercepted by an ANC. You may want to configure an ANC interception ACL for each ANC in an AppNav Cluster.

For information on how to configure an ANC interception ACL, see the "Configuring Interception Access Control Lists" section on page 5-28.

Configuring AppNav Cluster Settings

To configure AppNav Cluster settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **AppNav Clusters > All AppNav Clusters**.

The Manage AppNav Clusters window appears, which shows the status of each cluster.

From this window, you can perform the following tasks:

- View an AppNav Cluster topology and edit its settings by clicking on a cluster name.
- Delete an AppNav Cluster by selecting an AppNav Cluster and clicking the **Delete** icon in the taskbar of the Manage AppNav Clusters area.
- Create a new AppNav Cluster as described in the steps that follow.

Step 2 Click the name of the cluster whose settings you want to edit.

The cluster topology diagram appears.

Step 3 Choose **Configure > AppNav Cluster > AppNav Cluster**.

The Cluster Configuration window appears.

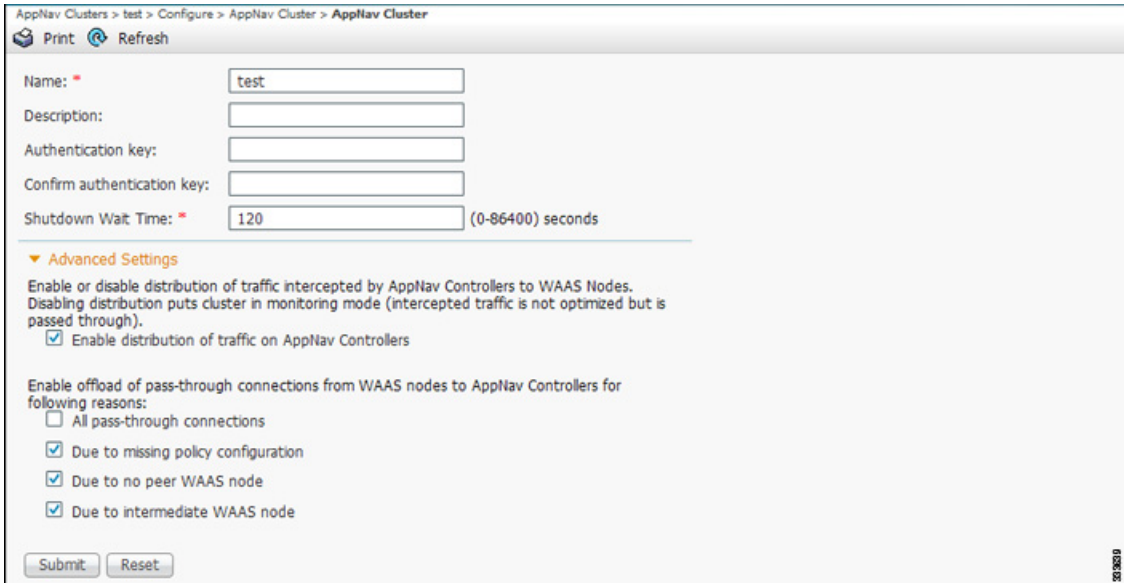

Step 4 In the Name field, enter a new name for the cluster if you want to rename it.

Step 5 (Optional) In the Description field, enter the cluster description. Use only letters and numbers, up to a maximum of 200 characters.

Step 6 (Optional) In the Authentication Key and Confirm Authentication Key fields, enter an authentication key that is used to authenticate communications between the WAAS devices in the cluster. Use only letters and numbers, up to a maximum of 64 characters.

Step 7 (Optional) In the Shutdown Wait Time field, enter the number of seconds that WNs in the cluster should wait for all connections to terminate before shutting down. The default is 120 seconds.

Step 8 (Optional) To configure cluster distribution and off-loading of pass-through connections, expand the **Advanced Settings** section by clicking it.

Step 9 (Optional) To enable distribution of traffic from the ANCs in the cluster to WNs, ensure that the **Enable distribution of traffic on AppNav Controllers** check box is checked. To disable distribution of traffic, uncheck this box. When distribution is disabled, the cluster operates in monitoring mode where it continues to intercept traffic and, instead of distributing it to WNs, passes it through. This mode can be useful for monitoring traffic statistics without optimizing the traffic.

Step 10 (Optional) To configure offloading of pass-through connections from WNs to ANCs, check the check boxes in the **Enable offload of pass-through connections from WAAS nodes to AppNav Controllers for following reasons** section. This feature allows pass-through connections to be passed through at the ANC instead of being distributed to the WN and then passed-through. Configure pass-through offload as follows:

a. To offload all pass-through connections, which includes connections passed through due to error conditions, check the **All pass-through connections** check box. Check this box only if you do not need application visibility on the WNs into pass-through traffic due to error conditions. The default is unchecked.

b. To offload connections passed through due to missing policy configuration, check the **Due to missing policy configuration** check box. The default is checked.

c. To offload connections passed through due to no peer WN, check the **Due to no peer WAAS node** check box. The default is checked.

d. To offload connections passed through due to an intermediate WN, check the **Due to intermediate WAAS node** check box. The default is checked.

e. If some of the WNs use different pass-through offload settings, you can synchronize the settings on all WNs to match the configuration shown here by checking the **Synchronize settings on all devices** check box. This check box is shown only if the settings on some WNs are different. The default is unchecked.

Step 11 Click **Submit**.

The lower part of this window shows lists of the ANCs, WNs, and WNGs that are part of the cluster. The controls in these parts of this window work as described in the following sections:

- AppNav Controllers—Configuring AppNav Controller Settings, page 4-28
- WAAS Nodes—Configuring WAAS Node Settings, page 4-29
- WAAS Node Groups—Configuring WAAS Node Group Settings, page 4-30

Configuring AppNav Controller Settings

To configure ANC settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **AppNav Clusters >** *cluster-name*.

Step 2 Click the **AppNav Controllers** tab below the topology diagram.

All ANCs in the cluster are listed, showing the name, location, IP address, interface used for intra-cluster traffic, and enabled status.

From this list, you can perform the following tasks:

- Edit the interface settings for an ANC by choosing the ANC and clicking the **Edit** taskbar icon.
- Delete an ANC by choosing the ANC and clicking the **Delete** taskbar icon.
- Add a new ANC to the cluster by clicking the **Add AppNav Controller** taskbar icon. See the "Adding an ANC to a Cluster" section on page 4-31.
- Enable a disabled ANC by choosing the cluster and clicking the **Enable** taskbar icon.
- Disable an ANC by choosing the ANC and clicking the **Disable** taskbar icon.

Step 3 Click the radio button next to the ANC that you want to edit and click the **Edit** taskbar icon.

The Edit AppNav Controller pane appears.

Step 4 If you want to enable optimization on the ANC, check the **Enable WAN optimization (Internal WAAS Node)** check box.

Step 5 If you enabled WAN optimization, from the **WAAS Node Group** drop-down list, choose the WNG to which the internal WN should belong.

Step 6 Click **Next**.

Step 7 (Optional) Configure the WCCP settings for the ANC. This screen does not appear if the ANC is configured for inline interception. For more information on the WCCP fields, see the "Configuring or Viewing the WCCP Settings on ANCs" section on page 5-22.

When finished with the WCCP settings, click **Next**. The graphical interface wizard appears.

Step 8 In the graphical interface view, configure interfaces on the AppNav Controller Interface Module as needed. For details on how to use the wizard, see the "Configuring Interfaces with the Graphical Interface Wizard" section on page 4-17.

Step 9 From the Cluster Interface drop-down list, select the interface to be used for intra-cluster traffic.

Step 10 (Optional) To enable swapping of client and WAAS device source IP address fields in intra-cluster traffic, check the **Enable swapping of source IP address in intra-cluster traffic** check box.

You may want to enable this option if you are using a port channel for the cluster interface or there is a load balancing device between the ANC and WN. This option may improve the load balancing of traffic that the ANC distributes to WNs for optimization because it load balances based on the client IP address rather than the ANC IP address. (For traffic from the server to the client, it swaps the server IP address with the ANC IP address.) The Central Manager enables this feature automatically if any existing ANCs have port channel cluster interfaces.

Step 11 Click **Finish**.

Configuring WAAS Node Settings

All WNs in the cluster must be configured with application-accelerator device mode and appnav-controller interception mode. If you created the cluster with the Central Manager AppNav Wizard, both of these settings are already done. (The wizard sets the interception mode and the device mode would have been set before running the wizard.)

From within the AppNav Cluster context, you can configure the following settings for a WN:

- WNG to which the WN belongs
- AppNav Controller Interface Module interface settings (including configuring port channel, standby, and bridge group interfaces)
- Choose the cluster interface used for intra-cluster traffic

To configure WN settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **AppNav Clusters >** *cluster-name*.

Step 2 Click the **WAAS Nodes** tab below the topology diagram.

All WNs in the cluster are listed, showing the name, location, IP address, interface in use, WNG to which the node belongs, and enabled status.

From this list, you can perform the following tasks:

- Edit the settings for a WN by choosing the WN and clicking the **Edit** taskbar icon.
- Delete a WN by choosing the WN and clicking the **Delete** taskbar icon.
- Add a new WN to the cluster by clicking the **Add WAAS Node** taskbar icon. See the "Adding a New WAAS Node to the Cluster" section on page 4-32.
- Enable a disabled WN by choosing the node and clicking the **Enable** taskbar icon.
- Disable a WN by choosing the node and clicking the **Disable** taskbar icon.

Step 3 Click the radio button next to the WN that you want to edit and click the **Edit** taskbar icon.

The WAAS Node pane appears.

Step 4 From the WAAS Node Group drop-down list, choose the WNG to which you want to assign the node.

Step 5 In the graphical interface view, configure interfaces on the AppNav Controller Interface Module as needed. For details on how to use the wizard, see the "Configuring Interfaces with the Graphical Interface Wizard" section on page 4-17.

Step 6 From the Cluster Interface drop-down list, select the interface to be used for intra-cluster traffic.

Step 7 (Optional) To enable swapping of client and WAAS device source IP address fields in intra-cluster traffic, check the **Enable swapping of source IP address in intra-cluster traffic** check box.

You may want to enable this option if you are using a port channel for the cluster interface or there is a load balancing device between the ANC and WN. This option may improve the load balancing of traffic that the ANC distributes to WNs for optimization because it load balances based on the client IP address rather than the ANC IP address. (For traffic from the server to the client, it swaps the server IP address with the ANC IP address.) The Central Manager enables this feature automatically if any existing ANCs have port channel cluster interfaces.

Step 8 Click **OK** to save the settings.

Configuring WAAS Node Group Settings

To configure WNG settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **AppNav Clusters >** *cluster-name*.

Step 2 Click the **WAAS Node Groups** tab below the topology diagram.

All WNGs in the cluster are listed, showing the name, description, and the WNs contained in the group.

From this list, you can perform the following tasks:

- Edit the settings for a WNG by choosing the WNG and clicking the **Edit** taskbar icon.
- Delete a WNG by choosing the WNG and clicking the **Delete** taskbar icon.
- Add a new WNG to the cluster by clicking the **Add WAAS Node Group** taskbar icon. See the "Adding a New WAAS Node Group to the Cluster" section on page 4-34.

Step 3 Click the radio button next to the WNG that you want to edit and click the **Edit** taskbar icon.

Step 4 (Optional) In the Description field, enter a description of the WNG.

Step 5 Click **Save** to save the settings.

Adding and Removing Devices from the AppNav Cluster

This section includes these topics:

Adding an ANC to a Cluster

To add a new ANC to an AppNav Cluster, follow these steps:

Step 1 Configure basic device and network settings on the new ANC, and ensure that the device mode is set to appnav-controller.

Step 2 From the WAAS Central Manager menu, choose **AppNav Clusters >** *cluster-name*.

Step 3 Click the **AppNav Controllers** tab below the topology diagram.

Step 4 Click the **Add AppNav Controller** taskbar icon.

The Add AppNav Controllers pane appears.

Step 5 Select one or more ANCs in the AppNav Controller device list by checking the check boxes next to the device names. You can use the filter settings in the taskbar to filter the device list.

If there are devices that are ineligible to join the cluster, you can click **Show Ineligible Devices** to see them and the reasons why they are ineligible. You can use the filter settings to filter the list.

Step 6 Click **Next**.

Step 7 Configure the interception method, policy, WCCP settings (if using WCCP interception), and interfaces for each ANC device you are adding:

a. From the Interception Method drop-down list, choose **WCCP** or **Inline**.

b. From the AppNav Policy-Map drop-down list, choose the AppNav policy to apply to the ANC.

c. (Optional) If you want to enable optimization on the ANC devices, check the **Enable WAN optimization (Internal WAAS Node)** check box.

d. (Optional) If you enabled WAN optimization, from the WAAS Node Group drop-down list, choose the WNG to which the internal WN should belong.

e. Click **Next**.

f. (Optional) If you chose WCCP interception, configure the WCCP settings on the WCCP settings pane that appears. For details on WCCP settings, see the "Configuring or Viewing the WCCP Settings on ANCs" section on page 5-22. Remember to check the **Enable WCCP Service** check box to enable WCCP.

g. If you configured WCCP settings, click **Next**.

h. Use the Cluster Interface Wizard graphical interface to configure the ANC interfaces. If you chose inline interception, you must configure a bridge group interface. For details on using this wizard, see the "Configuring Interfaces with the Graphical Interface Wizard" section on page 4-17.

i. From the Cluster Interface drop-down list, select the interface to be used for intra-cluster traffic.

j. (Optional) To enable swapping of client and WAAS device source IP address fields in intra-cluster traffic, check the **Enable swapping of source IP address in intra-cluster traffic** check box.

You may want to enable this option if you are using a port channel for the cluster interface or there is a load balancing device between the ANC and WN. This option may improve the load balancing of traffic that the ANC distributes to WNs for optimization because it load balances based on the client IP address rather than the ANC IP address. (For traffic from the server to the client, it swaps the server IP address with the ANC IP address.) The Central Manager enables this feature automatically if any existing ANCs have port channel cluster interfaces.

k. Click **Next** to save the settings and continue with the next ANC you are adding. If this is the last ANC being added, click **Finish**.

After a convergence waiting period of up to two minutes, the new ANCs are available in the cluster for traffic interception and distribution. Traffic interception on the new ANCs is prevented until the devices have fully joined the cluster. You can monitor the ANC status as described in the "Monitoring an AppNav Cluster" section on page 4-34.

Removing an ANC from a Cluster

To gracefully remove an ANC from an AppNav Cluster, follow these steps:

Step 1 Disable the traffic interception path on the ANC. For an inline ANC, shut down the in-path interfaces, and for an ANC using WCCP, disable WCCP.

Traffic previously routed to this ANC is rerouted to other ANCs in the cluster.

Step 2 Disable the ANC:

a. From the WAAS Central Manager menu, choose **AppNav Clusters >** *cluster-name*.

b. Click the **AppNav Controllers** tab below the topology diagram.

c. Click the radio button next to the ANC that you want to disable and then click the **Disable** taskbar icon.

The ANC is disabled and the service unreachable alarm is raised on the other ANCs in the cluster.

Step 3 (Optional) To permanently remove the ANC, click the radio button next to the ANC that you want to remove and then click the **Delete** taskbar icon.

This action removes the ANC from the ANCG on all other ANCs and clears the service unreachable alarm on the other ANCs. If the ANC is configured for WCCP interception, all WCCP settings on the device are removed. If the ANC is also configured as a WN, the WN is removed from the cluster.

Step 4 (Optional) Power down the ANC.

Adding a New WAAS Node to the Cluster

To add a new WAAS node (WN) to a cluster, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **AppNav Clusters >** *cluster-name*.

Step 2 Click the **WAAS Nodes** tab below the topology diagram.

Step 3 Click the **Add WAAS Node** taskbar icon.

The Add WAAS Nodes pane appears.

Step 4 Select one or more WNs in the WAAS Nodes device list by checking the check boxes next to the device names. You can use the filter settings in the taskbar to filter the device list.

If there are devices that are ineligible to join the cluster, click **Show Ineligible Devices** to see them and the reasons why they are ineligible. You can use the filter settings to filter the list.

Step 5 Click **Next**.

Step 6 Configure the WNG and interfaces for each WN device you are adding.

a. From the WAAS Node Group drop-down list, choose the WNG to which you want to add the new WNs. The list shows defined WNGs.

b. Click **Next**.

c. Use the Cluster Interface Wizard graphical interface to configure the WN interfaces. For details on using this wizard, see the "Configuring Interfaces with the Graphical Interface Wizard" section on page 4-17.

d. From the Cluster Interface drop-down list, select the interface to be used for intra-cluster traffic.

e. (Optional) To enable swapping of client and WAAS device source IP address fields in intra-cluster traffic, check the **Enable swapping of source IP address in intra-cluster traffic** check box.

   You may want to enable this option if you are using a port channel for the cluster interface or there is a load balancing device between the ANC and WN. This option may improve the load balancing of traffic that the ANC distributes to WNs for optimization because it load balances based on the client IP address rather than the ANC IP address. (For traffic from the server to the client, it swaps the server IP address with the ANC IP address.) The Central Manager enables this feature automatically if any existing ANCs have port channel cluster interfaces.

f. Click **Next** to save the settings and continue with the next WN you are adding. If this is the last WN being added, click **Finish**.

Step 7 Configure and enable optimization on the WNs. For details on configuring optimization, see Chapter 13, "Configuring Application Acceleration."

After a convergence waiting period of up to two minutes, the new WNs are available on all the ANCs for optimization.

Removing a WAAS Node from a Cluster

To remove a WAAS node (WN) from a cluster, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **AppNav Clusters >** *cluster-name*.

Step 2 Click the **WAAS Nodes** tab below the topology diagram.

Step 3 Choose the node and click the **Disable** taskbar icon.

This causes a graceful exit of the WN from the cluster, where the ANCs stop sending new flows to the WN but continue to distribute existing flows to it until the connection count reaches zero or the maximum shutdown wait time expires.

Note The default shutdown wait time is 120 seconds. You can configure it from the Shutdown Wait Time field in the AppNav Cluster tab.

Step 4 (Optional) When the graceful exit process on the WN is complete (all existing connections have terminated), remove the WN from the WNG on the ANCs by choosing the node and clicking the **Delete** taskbar icon.

You can monitor the node status in the topology diagram in the upper part of the window. The colored status light indicator on the device turns gray when the node is no longer processing connections.

Step 5 (Optional) Power down the WN.

Adding a New WAAS Node Group to the Cluster

To add a new WNG to a cluster, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **AppNav Clusters >** *cluster-name*.

Step 2 Click the **WAAS Node Groups** tab below the topology diagram.

Step 3 Click the **Add WAAS Node Group** taskbar icon.

The Add WAAS Node Group pane appears.

Step 4 In the Name field, enter the name of the WNG.

Step 5 (Optional) In the Description field, enter a description of the WNG.

Step 6 Click **OK** to save the settings.

Step 7 Add one or more WNs to the new WNG. To add a new WN, see the "Adding a New WAAS Node to the Cluster" section on page 4-32, or to reassign an existing WN to the new WNG, see the "Configuring WAAS Node Settings" section on page 4-29.

After a convergence waiting period of up to two minutes, the new WNG is available on all the ANCs for optimization.

Removing a WAAS Node Group from a Cluster

To remove a WAAS node group (WNG) from a cluster, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **AppNav Clusters >** *cluster-name*.

Step 2 Click the **WAAS Nodes** tab below the topology diagram.

Step 3 For each WN in the WNG, click the radio button next to the node name and click the **Disable** taskbar icon. This causes a graceful exit of each WN from the cluster.

Step 4 After all WNs have completed a graceful exit from the cluster, click the **WAAS Node Groups** tab.

You can monitor the node status in the topology diagram in the upper part of the window. The colored status light indicator on a device turns gray when the node is no longer processing connections.

Step 5 (Optional) Choose the WNG you want to remove and click the **Delete** taskbar icon.

Monitoring an AppNav Cluster

To monitor an AppNav Cluster, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **AppNav Clusters >** *cluster-name*.

The cluster home window displays the cluster topology and device status (see Figure 4-5).

Figure 4-5 AppNav Cluster Topology and Status

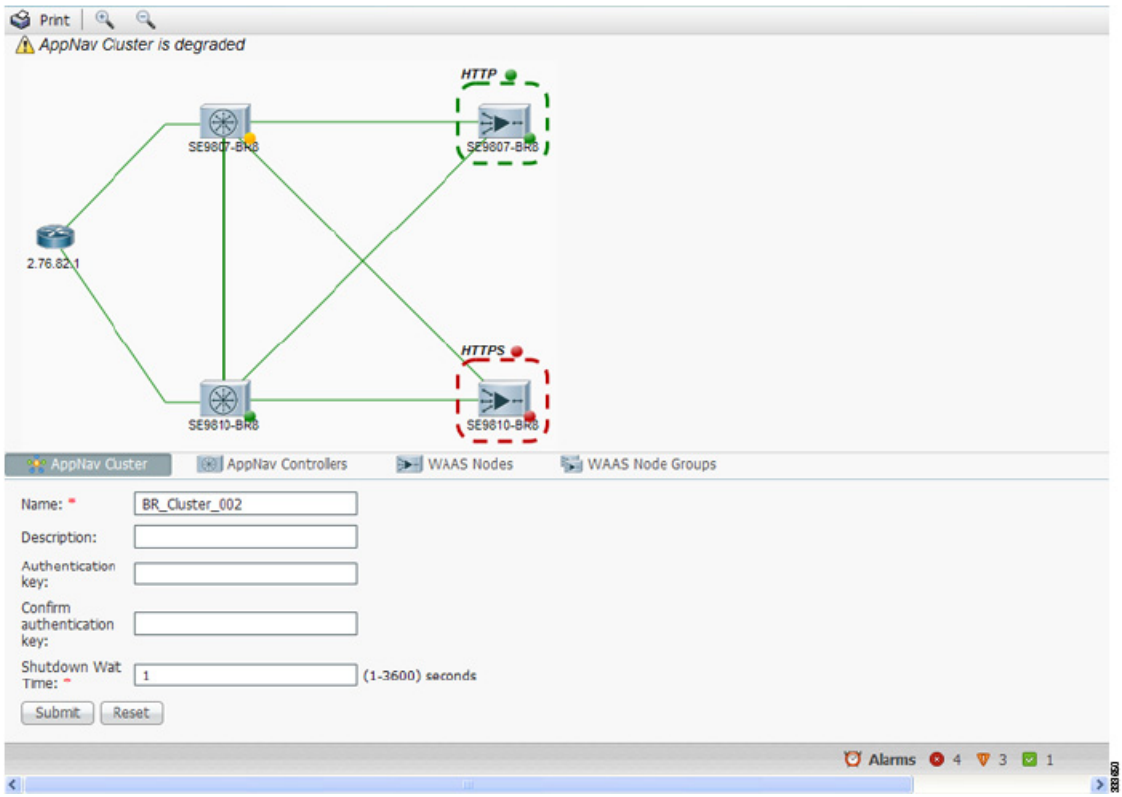

To zoom in or out on the topology diagram, click the + or – magnifying glass icons in the taskbar. You can also click on the diagram and drag it within the window to reposition it.

To change the cluster settings, edit any of the fields below the topology diagram and click **Submit**.

To see all ANCs, click the **AppNav Controllers** tab below the diagram. From this tab, you can edit, delete, add, enable, or disable an ANC in the cluster.

To see all WNs, click the **WAAS Nodes** tab below the diagram. From this tab, you can edit, delete, add, enable, or disable a WN in the cluster.

To see all WNGs, click the **WAAS Node Groups** tab below the diagram. From this tab, you can edit, delete, or add a WNG in the cluster.

The overall cluster status is shown in the top left corner of the diagram, as follows:

- Green—All ANCs are operational with no error conditions.
- Yellow—Degraded because one or more ANCs have operational issues. This is also the initial state before all nodes have sent status updates.
- Red—Cluster is down because all ANCs are down or indicates a split cluster where there is no connectivity between one or more ANCs.

The overall cluster status does not include administratively disabled ANCs.

The colored status light indicators on each device and dotted lines around each WNG show the status of the device or group:

- Green—Operational with no error conditions

- Yellow—Degraded (overloaded, joining cluster, or has other noncritical operational issues)
- Red—Critical (one or more processes is in a critical state)
- Gray—Disabled
- Black—Unknown status

The colored lines between each device show the status of the link between devices:

- Green—Operational with no error conditions
- Red—Link is down
- Black—Unknown status

An orange triangle warning indicator is shown on any device for which the Central Manager may not have current information because the device has not responded within the last 30 seconds (the device could be offline or unreachable).

Note A recently removed device still appears in the topology diagram for a few minutes until all devices agree on the new cluster topology.

To view a more comprehensive device status display, hover your cursor over a device icon to see the 360-degree Network Device View popup window (Figure 4-6). The popup window for a WN device is similar.

Figure 4-6 ANC 360-Degree Network Device View

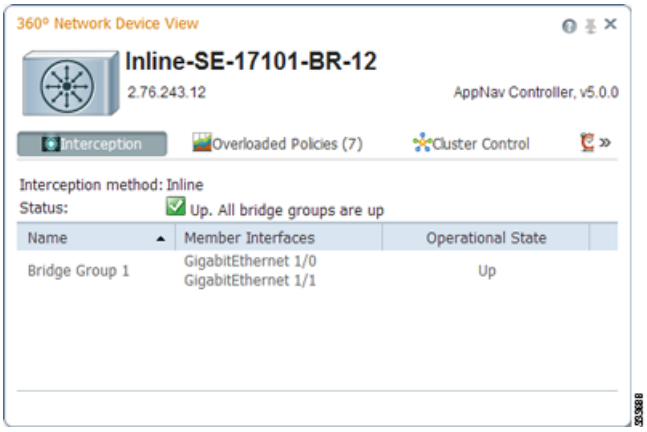

The 360-degree Network Device View shows the following status information:

- Device name and IP address
- Device type and software version
- (ANC only) Interception tab that displays the interception method: Inline or WCCP. For inline, this tab shows the bridge groups defined for interception, their member interfaces, and their status. For WCCP, this tab lists the defined WCCP service IDs, their associated client IP addresses, router IP address, and notes about problems.
- (ANC only) Overloaded Policies tab that lists any monitored AppNav policies that are overloaded.
- (ANC only) Cluster Control tab that lists all devices in the cluster, with device name, IP address, service type, liveliness state, and reason for any error condition

- (WN only) Optimization tab that lists the application accelerators and their status
- Alarms tab that lists pending alarms on the device
- Interfaces tab that lists the device interfaces and status. You can filter the list by choosing a filter type from the drop-down list above the interface list, entering filter criteria, and clicking the filter icon.

You can pin the status popup window so it stays open by clicking the pin icon in the upper right corner. You can also drag the popup to any location within your browser window.

For additional cluster status, you can view the Monitor > AppNav > AppNav Report as described in the "AppNav Report" section on page 17-43.

If you have multiple AppNav Clusters, you can see brief status for all at once by choosing **AppNav Clusters > All AppNav Clusters** from the menu.

To trace connections, see the "AppNav Connection Tracing" section on page 4-37.

Note You may see a taskbar icon named Force Settings on all Devices in a Group if the configuration across all ANCs in the cluster becomes unsynchronized. If you see the icon, the cluster settings, ANC configuration, WN configuration, and WNG configuration do not match on all ANCs in the cluster. This problem can occur if you configure a device outside the Central Manager by using the CLI. Click this taskbar icon to update all devices with the configuration that is currently shown in the Central Manager for the cluster.

AppNav Connection Tracing

To assist in troubleshooting AppNav flows, you can use the Connection Trace tool in the Central Manager. This tool shows the following information for a particular connection:

- If the connection was passed through or distributed to a WNG
- Pass-through reason, if applicable
- The WNG and WN to which the connection was distributed
- Accelerator monitored for the connection
- Class-map applied

To use the Connection Trace tool, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **AppNav Clusters >** *cluster-name*.

Step 2 Choose **Monitor > Tools > Connection Trace**.

Step 3 In the AppNav Controller drop-down list, choose the ANC that has the connection you want to trace.

Step 4 From the Site (Remote Device) drop-down list, choose the peer WAAS device at the remote site.

Step 5 In one or more of the Source IP, Source Port, Destination IP, and Destination Port fields, enter matching criteria for one or more connections.

Step 6 Click **Trace** to display the connections that match the IP address and port criteria.

Connections are displayed in the Connection Tracing Results table below the fields. Use the filter settings in the Show drop-down list to filter the connections as needed. You can use a quick filter to filter on any value or show all connections.

You can display flow distribution information from the CLI by using the **show appnav-controller flow-distribution** EXEC command.

Another troubleshooting tool that you can use to trace connections is the WAAS Tcptraceroute tool. For details, see the "Using WAAS TCP Traceroute" section on page 17-61.

CHAPTER 5

Configuring Traffic Interception

This chapter describes how to configure interception of TCP traffic in an IP-based network, based on the IP and TCP header information and how to redirect the traffic to WAAS devices. This chapter describes the use of the Web Cache Communication Protocol (WCCP), policy-based routing (PBR), inline mode for transparent redirection of traffic to WAEs, appnav-controller mode for use with an AppNav Controller, and VPATH interception for redirection of VMware packets to virtual WAAS (vWAAS).

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE and WAVE appliances, WAE Network Modules (the NME-WAE family of devices), SM-SRE modules running WAAS, and vWAAS instances.

Before you do the procedures in this chapter, you should complete a basic initial installation and configuration of your WAAS network as described in the *Cisco Wide Area Application Services Quick Configuration Guide*. For detailed command syntax information for any of the CLI commands in this chapter, see the *Cisco Wide Area Application Services Command Reference*. For more information about WCCP see the CISCO IOS documentation.

This chapter contains the following sections:

Information About Interception Methods

In a WAAS network, traffic between clients in the branch offices and the servers in the data center can be redirected to WAEs for optimization, redundancy elimination, and compression. Traffic is transparently intercepted and redirected to WAEs based on policies that have been configured on the routers or on an AppNav Controller (ANC). The network elements that transparently redirect requests

to a local WAE can be a router using WCCP Version 2 or PBR to redirect traffic to the local WAE or a Layer 4 to Layer 7 switch (for example, the Catalyst 6500 series Content Switching Module [CSM] or Application Control Engine [ACE]). Alternately, you can intercept traffic directly by using the inline mode with a WAE that has a Cisco WAE Inline Network Adapter or Interface Module. When equipped with a Cisco AppNav Controller Interface Module, a WAVE appliance or cluster can intercept network traffic through WCCP or inline mode and, based on flow policies, distribute that traffic to one or more WAEs (WAAS nodes) for optimization.

Table 5-1 summarizes the transparent traffic interception methods that are supported in your WAAS network.

Table 5-1 Supported Methods of Transparent Traffic Interception

Method	Comment
WCCP Version 2	Used for transparent interception of application traffic and Common Internet File System (CIFS) traffic. Used in branch offices and data centers to transparently redirect traffic to the WAAS devices. The traffic is transparently intercepted and redirected to the local WAE or ANC by a WCCP-enabled router or a Layer 3 switch. You must configure WCCP on the router and WAE in the branch office and the router and WAE in the data center. For more information, see the following sections: • Information About WCCP Interception, page 5-3 • Configuring Advanced WCCP Features on Routers, page 5-6 • Configuring WCCP on WAEs, page 5-11
PBR	In branch offices, used for wide area application optimization. The branch office router is configured to use PBR to transparently intercept and route both client and server traffic to the WAE that resides in the same branch office. In data centers, used for data center application optimization. The data center router or Layer 3 switch may be configured to use PBR to transparently intercept and route client and server traffic to WAEs within the data center. PBR, however, does not support load balancing across multiple WAEs (such as WCCP does). Neither does it support load balancing when you are using a hardware load balancer, such as the Cisco CSM or ACE. See the "Using Policy-Based Routing Interception" section on page 5-33.
Inline	The WAE physically and transparently intercepts traffic between the clients and the router. To use this mode, you must use a WAAS device with the Cisco WAE Inline Network Adapter, Cisco Interface Module, or Cisco AppNav Controller Interface Module installed. See the "Using Inline Mode Interception" section on page 5-42.
VPATH	Used for VPATH interception on vWAAS devices. See the "Configuring VPATH Interception on a vWAAS Device" section on page 5-55.
AppNav Controller	For WAEs that are part of an AppNav deployment and are configured as WAAS nodes in an AppNav Cluster, you must configure them to use the appnav-controller interception method. This configuration allows WAEs to receive and optimize traffic that is intercepted and distributed by the AppNav Controllers. See the "Configuring AppNav Interception" section on page 5-56.
ACE or CSM	Cisco Application Control Engine (ACE) or Catalyst 6500 series Content Switching Module (CSM) installed in the data center for data center application optimization. The ACE or CSM allows for both traffic interception and load balancing across multiple WAEs within the data center.

If a WAE device is behind a firewall that prevents traffic optimization, you can use the directed mode of communicating between peer WAEs over the WAN. For details, see the "Configuring Directed Mode" section on page 6-27.

Information About WCCP Interception

The WAAS software uses the WCCP standard, Version 2 for redirection. The main features of WCCP Version 2 include support for the following:

- Up to 32 WAEs per WCCP service
- Up to 32 routers per WCCP service
- Authentication of protocol packets
- Redirection of non-HTTP traffic
- Packet return (including generic routing encapsulation [GRE], allowing a WAE to reject a redirected packet and to return it to the router to be forwarded)
- Masking for improved load balancing
- Multiple forwarding methods
- Packet distribution method negotiation within a service group
- Command and status interaction between the WAE and a service group

Note WCCP works only with IPv4 networks.

WAAS software supports the WCCP TCP promiscuous mode service (services 61 and 62 by default, though these service IDs are configurable). This WCCP service requires that WCCP Version 2 is running on the router and the WAE.

The TCP promiscuous mode service is a WCCP service that intercepts all TCP traffic and redirects it to the local WAE.

The WAAS software also supports service passwords, WAE failover, flow protection, and interception ACLs.

Many Cisco routers and switches can be configured and enabled with WCCP Version 2 support for use with WAAS devices.

Note Many legacy Cisco routers, including the 2500, 2600, and 3600 routers, have far less processing power and memory than newer routing platforms such as the Integrated Services Router (ISR) models 2800 and 3800. As such, the use of WCCPv2 or PBR may cause a high level of CPU utilization on the router and cause erratic behavior. WAAS can be configured to work with these routers, but not to the same levels of performance or scalability as can be found with newer routing platforms. The Cisco ISR is the routing platform of choice for the branch office.

If you are experiencing erratic behavior, such as the WAE being ejected from the service group, enable fair-queuing, weighted fair-queuing, or rate-limiting on all physical interfaces on the router that connect to users, servers, WAEs, and the WAN. Fair-queuing cannot be configured on subinterfaces, and should be configured on both ingress and egress physical interfaces. If another form of queuing is already configured on the LAN or WAN interfaces other than fair-queuing that provides similar fairness, it should be sufficient.

Additionally, limit the amount of bandwidth that can be received on the LAN-side interface of the router, to help the router keep its interface queues less congested and provide better performance and lower CPU utilization. Set the maximum interface bandwidth on the router to no more than 10 times the WAN bandwidth capacity. For instance, if the WAN link is a T1, the LAN interface and WAE LAN interface bandwidth should be throttled to 10 * T1 = 10 * 1.544 Mbps, or approximately 15 Mbps. See the Cisco IOS documentation for more information.

This section contains the following topics:

Guidelines for Configuring WCCP

When you configure transparent redirection on a WAE using WCCP Version 2, follow these guidelines:

- Intercept and redirect packets on the inbound interface whenever possible.
- Use WCCP GRE or generic GRE as the egress method if you want to place WAEs on the same VLAN or subnet as clients and servers. This topology is not allowed when using the IP forwarding egress method.
- Branch WAEs must not have their packets encrypted or compressed and should be part of the "inside" Network Address Translation (NAT) firewall if one is present.
- Use Layer 2 redirection as the packet forwarding method if you are using Catalyst 6500 series switches or Cisco 7600 series routers. Use Layer 3 GRE packet redirection if you are using any other Cisco series router.
- When you configure WCCP for use with the Hot Standby Router Protocol (HSRP), you must configure the WAE with the HSRP or the Virtual Router Redundancy Protocol (VRRP) virtual router address as its default gateway, and the WAE WCCP router-list with the primary address of the routers in the HSRP group.
- CEF is required for WCCP and must be enabled on the router.
- Place branch WAEs on the client side of the network to minimize client-side packets through the router.
- Use WCCP passwords to avoid denial-of-service attacks. For more information, see the "Setting a Service Group Password on a Router" section on page 5-10.
- Use WCCP redirect lists for new implementations to limit client or server populations. For more information, see the "Configuring IP Access Lists on a Router" section on page 5-9.
- You must configure the WAE to accept redirected packets from one or more WCCP-enabled routers.
- To configure basic WCCP, you must enable the WCCP service on at least one router in your network and on the WAE or ANC that you want the traffic redirected to. It is not necessary to configure all of the available WCCP features or services to get your WAE up and running. For an example of how to complete a basic WCCP configuration on routers and WAEs in a branch office and data center, see the *Cisco Wide Area Application Services Quick Configuration Guide*.
- You must configure the routers and WAEs to use WCCP Version 2 instead of WCCP Version 1 because WCCP Version 1 only supports web traffic (port 80).

- After enabling WCCP on the router, you must configure the TCP promiscuous mode service on the router and the WAE, as described in the *Cisco Wide Area Application Services Quick Configuration Guide*. The service IDs are configurable on the WAE and you can choose a pair of numbers different from the default of 61 and 62 to allow the router to support multiple WCCP farms because the WAEs in different farms can use different service IDs. The router configuration must use WCCP service IDs that match those configured on the WAEs in each farm that it is supporting.
- In order for the WAE to function in TCP promiscuous mode, the WAE uses WCCP Version 2 services 61 and 62 (the service IDs are configurable). These two WCCP services are represented by the canonical name tcp-promiscuous on the WAE.
- You can use CLI commands to configure basic WCCP on both the routers and the WAEs or ANCs, or you can use CLI commands to configure the router for WCCP and use the WAAS Central Manager to configure basic WCCP on the WAEs or ANCs. In the configuration example provided in the *Cisco Wide Area Application Services Quick Configuration Guide*, the **wccp** global configuration command is used to configure basic WCCP on the WAEs or ANCs.

  We recommend that you use the WAAS CLI to complete the initial basic configuration of WCCP on your first branch WAE and data center WAE, as described in the *Cisco Wide Area Application Services Quick Configuration Guide.* After you have verified that WCCP transparent redirection is working properly, you can use the WAAS Central Manager to modify this basic WCCP configuration or configure additional WCCP settings (for example, load balancing) for a WAE. For more information, see the "Configuring WCCP on WAEs" section on page 5-11. After you have configured basic WCCP on the router, you can configure advanced WCCP features on the router, as described in the "Configuring Advanced WCCP Features on Routers" section on page 5-6.
- To ensure consistency among WAEs, we recommend that you configure WCCP settings on one device and then use the **Copy Settings** taskbar icon from within the WCCP configuration window to copy the settings to other devices in your network. You should copy the settings only to WAEs in the same WCCP service farm, AppNav Controller group (ANCG), or WAAS node group (WNG), since WCCP settings may need to be different in different farms or service groups.
- When you add a new router to an existing WCCP router farm or WCCP service group, the new router will reset existing connections. Until WCCP reestablishes path redirections and assignments, packets are sent directly to the client (as expected).
- The router must support the redirect and return methods configured on the WAE. If the router does not support the configured methods, the WAE will not join the WCCP router farm. If you have a mix of routers in the farm, only those routers that support the configured methods will join the farm.
- The WAE only joins the WCCP farm if the assignment method configured on the WAE is supported by the router. (The strict assignment method is always enforced with version 4.4.1 and later.)
- A WAE joins a WCCP farm only if it is seen by all the configured routers in the farm. If there is a link failure with any one of the routers, the farm reconfigures and the WAE is removed from the farm.
- All WAEs in a WCCP farm must use the same pair of WCCP service IDs (the default is 61 and 62), and these IDs must match all routers that are supporting the farm. A WAE with different WCCP service IDs is not allowed to join the farm and an alarm is raised. Likewise, all WAEs in a farm must use the same value for the failure detection timeout. A WAE raises an alarm if you configure it with a mismatching value.
- VPN routing and forwarding (VRF)-aware WCCP scalability is as follows:
  - The maximum number of WAEs supported by a single VRF instance is 32.
  - The maximum number of VRF instances supported by the router is router dependent.
  - VRF-aware WCCP is supported only on specific releases of Cisco IOS software. Ensure that the router is running a release of Cisco IOS software that supports VRF-aware WCCP.

  - Each VRF instance has independent assignment, redirection, and return methods.
- In a WAAS AppNav deployment, enable WCCP only on the ANC devices that are intercepting traffic and distributing it to the optimizing WAAS nodes (WNs). Configure WNs that are part of the AppNav Cluster with the appnav-controller interception method.

Guidelines for File Server Access Methods

Some file servers have several network interfaces and can be reached through multiple IP addresses. For these server types, you must add all the available IP addresses to the branch WAE's WCCP accept list. This situation prevents a client from bypassing the branch WAE by using an unregistered IP address. The WAE Device Manager GUI displays all the IP addresses in the GUI.

Some file servers have several NetBIOS names and only one IP address. For these servers, if the client connects using the IP address in the UNC path (that is, \\IP_address\share instead of \\server\share), WAAS selects the first NetBIOS name from the server list in the WAE Device Manager GUI that matches this IP address. WAAS uses that name to perform NetBIOS negotiations between the data center WAE and the file server, and to create resources in the cache. If a file server uses multiple NetBIOS names to represent virtual servers (possibly with different configurations) and has one NetBIOS name that is identified as the primary server name, put that name in the server list before the other names.

Configuring Advanced WCCP Features on Routers

This section describes how to configure the advanced WCCP Version 2 features on a WCCP-enabled router that is transparently redirecting requests to WAEs in your WAAS network and contains the following topics:

Note Before you do the procedures in this section, you should have already configured your router for basic WCCP as described in the *Cisco Wide Area Application Services Quick Configuration Guide*.

Information About Configuring a Router to Support WCCP Service Groups

WCCP Version 2 enables a set of branch WAEs in a WAE or ANC group to connect to multiple routers. The WAEs in a group and the WCCP Version 2-enabled routers connected to the WAE group that are running the same WCCP service are known as a *service group*.

Through communication with the branch WAEs, the WCCP Version 2-enabled routers are aware of the available branch WAEs. Routers and branch WAEs become aware of one another and form a service group using WCCP Version 2. See Figure 5-1.

In a WAAS AppNav deployment, only the ANCs are included in the service group. The routers do not send traffic directly to the optimizing WAEs (WNs); instead, ANCs distribute traffic within the WAAS network to the optimizing WNs.

Figure 5-1 Service Groups with WCCP Version 2

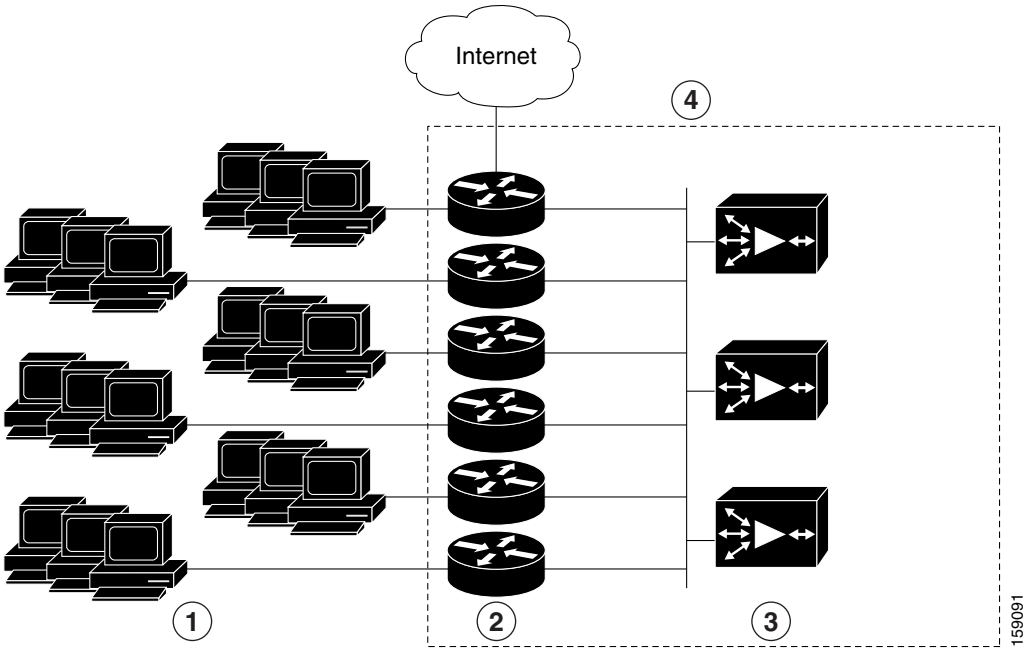

1	Clients requesting file services	**3**	Branch WAEs
2	Cisco routers	**4**	WAE service group

If you have a group of branch WAEs, the WAE that is seen by all the WCCP Version 2-enabled routers and that has the lowest IP address becomes the lead branch WAE.

The following procedure describes how a branch WAE in a service group is designated as the lead:

1. Each branch WAE is configured with a list of WCCP-enabled routers.

   Multiple WCCP-enabled routers can service a group (up to 32 routers can be specified). Any of the available routers in a service group can redirect packets to each of the branch WAEs in the group.

2. Each branch WAE announces its presence to each router on the router list. The routers reply with their view of branch WAEs in the service group.
3. After the view is consistent across all of the branch WAEs in the group, one branch WAE is designated as the lead branch WAE and sets the policy that the WCCP-enabled routers need to deploy in redirecting packets.

The lead branch WAE determines how traffic should be allocated across the branch WAEs in the group. The assignment information is passed to the entire service group from the designated lead branch WAE so that the WCCP-enabled routers of the group can redirect the packets and the branch WAEs in the group can better manage their load.

WCCP uses service groups to define WAAS services for a WCCP Version 2-enabled router and branch WAEs in a group. WCCP also redirects client requests to these groups in real time.

All ports receiving redirected traffic that are configured as members of the same WCCP service group share the following characteristics:

- They have the same hash or mask parameters, as configured with the WAAS Central Manager (the "Configuring or Viewing the WCCP Settings on WAEs" section on page 5-17) or the WAAS CLI (the **wccp service-number mask** global configuration command).
- The WCCP Version 2 service on individual ports cannot be stopped or started individually (a WCCP Version 2 restriction).

Configuring a Router to Support WCCP Service Groups

To direct a WCCP Version 2-enabled router to enable or disable support for a WCCP service group, use the **ip wccp** global configuration command. To remove the ability of a router to control support for a WCCP service group, use the **no** form of this command.

The following example shows how to enable the TCP promiscuous mode service (WCCP Version 2 services 61 and 62) on a router:

```
Router(config)# ip wccp 61
Router(config)# ip wccp 62
```

On each WAE, configure multiple router addresses in the WCCP router list, one for each router in the service group.

```
WAE(config)# wccp router-list 1 10.10.10.20 10.10.10.21
```

Finally, you need to configure each router for WCCP interception on the inbound direction of the appropriate interfaces, using commands similar to the following:

```
Router(config)# interface fa1/0.40
Router(config-subif)# ip wccp 61 redirect in
Router(config-subif)# exit
Router(config)# interface serial0
Router(config-subif)# ip wccp 62 redirect in
Router(config-subif)# exit
```

When a new WAE is brought online, it joins the WCCP service group. With a new WAE in the service group, the hash tables responsible for distributing the load are changed, and traffic that previously went to WAE1 may now go to WAE2. Flow protection must be enabled in order for WAE2 to forward packets of already connected clients to WAE1. The end result is that all requests that belong to a single session are processed by the same WAE. Without flow protection enabled, adding a WAE to the service group might disconnect some of the existing clients.

When an WAE is removed from the service group, its clients are disconnected (if they reconnect, they will reach another WAE, if one is available, or the origin file server).

WAAS supports WAE failover by reconnecting clients with other branch WAEs if a branch WAE crashed. In the event of a crash, the branch WAE stops issuing WCCP keepalives (constant high CPU load may also result in loss of keepalives and can also be considered a failover case). The router detects the lack of keepalives and removes the branch WAE from the service group. The designated branch WAE updates the WCCP configuration hash table to reflect the loss of the branch WAE and divides its buckets among the remaining branch WAEs. A new designated lead branch WAE is elected if the crashed one was the lead branch WAE. The client is disconnected, but subsequent connections are processed by another branch WAE.

Once a TCP flow has been intercepted and received by a branch WAE, the failure behavior is identical to that exhibited during nontransparent mode. For example, data center WAE and file server failure scenarios are not handled any differently as a result of using WCCP interception.

Note When you add a new router to an existing WCCP router farm or WCCP service group, the new router will reset existing connections. Until WCCP reestablishes path redirections and assignments, packets are sent directly to the client (as expected).

Configuring IP Access Lists on a Router

You can optionally configure the router to redirect traffic from your WAE based on access control lists (ACLs) that you define on the router. These access lists are also referred to as redirect lists.

Note We recommend that you use redirect lists on the WCCP-enabled router where possible, because that is the most efficient method to control traffic interception. However, you can also configure static bypass lists or interception ACLs on the WAEs, and of these two, we recommend using interception ACLs because they are more flexible and give better statistics about passed-through connections. For information about how to configure an interception ACL for a WAE, see the "Configuring Interception Access Control Lists" section on page 5-28. For information about how to configure a static bypass list, see the "Configuring Static Bypass Lists for WAEs" section on page 5-27. You can also configure interface ACLs on WAEs to control management access to the WAE, as described in Chapter 9, "Creating and Managing IP Access Control Lists for WAAS Devices."

Redirect lists that are configured on the routers have the highest priority, followed by static bypass lists or interception ACLs on WAEs. Interception ACLs that are configured on WAEs take precedence over any application definition policies that have been defined on the WAE.

A WCCP Version 2-enabled router can be configured with access lists to permit or deny redirection of TCP traffic to a WAE. The following example shows that traffic conforming to the following criteria are not redirected by the router to the WAE:

- Originating from the host 10.1.1.1 destined for any other host
- Originating from any host destined for the host 10.255.1.1

```
Router(config)# ip wccp 61 redirect-list 120
Router(config)# ip wccp 62 redirect-list 120
Router(config)# access-list 120 deny ip host 10.1.1.1 any
Router(config)# access-list 120 deny ip any host 10.1.1.1
Router(config)# access-list 120 deny ip any host 10.255.1.1
Router(config)# access-list 120 deny ip host 10.255.1.1 any
Router(config)# access-list 120 permit ip any
```

Traffic not explicitly permitted is implicitly denied redirection. The **access-list 120 permit ip any** command explicitly permits all traffic (from any source on the way to any destination) to be redirected to the WAE. Because criteria matching occurs in the order in which the commands are entered, the global **permit** command is the last command entered.

To limit the redirection of packets to those packets matching an access list, use the **ip wccp redirect-list** global configuration command. Use this command to specify which packets should be redirected to the WAE.

When WCCP is enabled but the **ip wccp redirect-list** command is not used, all packets matching the criteria of a WCCP service are redirected to the WAE. When you specify the **ip wccp redirect-list** command, only packets that match the access list are redirected.

The **ip wccp** global configuration command and the **ip wccp redirect** interface configuration command are the only commands required to start redirecting requests to the WAE using WCCP. To instruct an interface on the WCCP-enabled router to check for appropriate outgoing packets and redirect them to a WAE, use the **ip wccp redirect** interface configuration command. If the **ip wccp** command is enabled but the **ip wccp redirect** command is disabled, the WCCP-enabled router is aware of the WAE but does not use it.

To specify the access list by name or number, use the **ip wccp group-list** global configuration command, which defines criteria for group membership. In the following example, the **access-list 1 permit 10.10.10.1** command is used to define the IP address of the WAE that is allowed to join the WCCP service group:

```
Router(config)# ip wccp 61 group-list 1
Router(config)# ip wccp 62 group-list 1
Router(config)# access-list 1 permit 10.10.10.1
```

Tip If you have a WCCP service farm with multiple WAEs, the load balancing assignment may cause packets that are sent to the WAE devices themselves (such as management traffic) to be redirected to a different WAE in the farm, negatively impacting performance. To avoid this situation, we recommend that you configure a WCCP redirect list that excludes traffic that is sent to the WAE IP addresses from being redirected.

For more information on access lists, see the Cisco IOS IP addressing and services documentation.

Setting a Service Group Password on a Router

For security purposes, you can set a service password for your WCCP Version 2-enabled router and the WAEs that access it. Only devices configured with the correct password are allowed to participate in the WCCP service group.

From the global configuration mode of your WCCP-enabled router, enter the following commands to specify the service group password for the TCP promiscuous mode service on the router (the service IDs must match the service IDs configured on the WAE):

```
Router(config)# ip wccp 61 password [0-7] password
Router(config)# ip wccp 62 password [0-7] password
```

The required *password* argument is the string that directs the WCCP Version 2-enabled router to apply MD5 authentication to messages received from the specified service group. Messages that are not accepted by the authentication are discarded. *0-7* is the optional value that indicates the HMAC MD5 algorithm used to encrypt the password. This value is generated when an encrypted password is created for the WAE. 7 is the recommended value. The optional *password* argument is the optional password name that is combined with the HMAC MD5 value to create security for the connection between the router and the WAE.

For information about how to use the WAAS Central Manager to specify the service group password on a WAE, see the "Configuring or Viewing the WCCP Settings on WAEs" section on page 5-17.

Configuring a Loopback Interface on the Router

The highest IP address among the router's loopback interfaces is used to identify the router to the WAEs.

The following example configures the loopback interface, exits configuration mode, and saves the running configuration to the startup configuration:

```
Router(config)# interface Loopback0
Router(config-if)# ip address 111.111.111.111 255.255.255.0
Router(config-if)# no shutdown
Router(config-if)# end
Router# copy running-config startup-config
```

Configuring Router QoS for WCCP Control Packets

WAAS sends WCCP control packets marked with a differentiated services code point (DSCP) value of 192. (In WAAS versions earlier than 4.2, packets were unmarked.) For a router to honor this priority value, you must configure the router's multilayer switching (MLS) quality of service (QoS) port trust state and classify traffic by examining the DSCP value. To configure the router appropriately, use the **mls qos trust dscp** command in interface configuration mode on the interface connected to the WAE.

Configuring WCCP on WAEs

This section contains the following topics:

Note Before you do the procedures in this section, you should have completed an initial configuration of your WAAS network, which includes the basic configuration of WCCP Version 2 and the TCP promiscuous mode service on your routers and WAEs, as described in the *Cisco Wide Area Application Services Quick Configuration Guide*.

Information About Load Balancing and WAEs

Multiple WAEs with WCCP support can be deployed for dynamic load balancing to enable adjustments to the loads being forwarded to the individual WAEs in a service group. IP packets received by a WCCP-enabled router are examined to determine if it is a request that should be directed to a WAE. Packet examination involves matching the request to a defined service criteria. These packets are passed to the processing routine on the router to determine which WAE, if any, should receive the redirected packets.

Note In a WAAS AppNav deployment, only the ANCs are included in the service group and are load balanced by the routers. The routers do not send traffic to the optimizing WAEs (WNGs); instead, ANCs distribute traffic to the optimizing WNGs.

You can use load balancing to balance the traffic load across multiple WAEs. Load balancing allows the set of hash address buckets assigned to a WAE to be adjusted, shifting the load from an overwhelmed WAE to other WAEs that have available capacity. Two assignment methods are used by this technique: hashing and masking.

Assignment method denotes the method used by WCCP to perform load distribution across WAEs. The two possible load-balancing assignment methods are hashing and masking. If the mask load-balancing method is not specified, then the hash load-balancing method, which is the default method, is used.

Note In a WAAS AppNav deployment, only the mask assignment method is supported and is the default.

WCCP supports redirection based on a hash function. The hash key may be based on the source or destination IP address of the packet. For WAAS, load-balancing hashing is based on a source IP address (default), a destination IP address, or both.

The hash function uses the source IP address to obtain an address bucket to which the packet is assigned. These source address buckets are then mapped to a particular WAE depending on how many WAEs are present and how busy they are. (See Figure 5-2.)

Figure 5-2 Load Balancing Through Hashing of IP Addresses

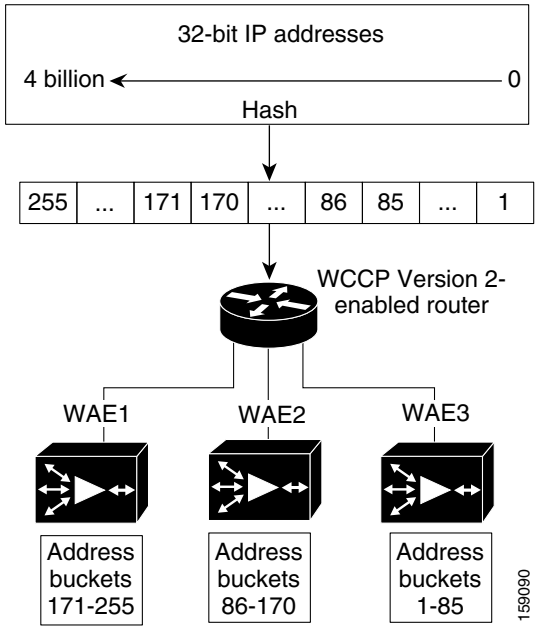

Note Packets that the WAEs do not service are tunneled back to the same router from which they were received. When a router receives a formerly redirected packet, it knows not to redirect it again.

Destination IP address hashing guarantees that a single WAE caches a given file server. This method, which allows a local coherency directive to be safely applied to the file server content (provided that no other collaboration on the content occurs), improves performance and WAN link and disk utilization. This method may distribute the load unevenly because of uneven activity on a file server.

Source IP address hashing has better potential for session distribution between the caches on branch WAEs. This method may impact performance and WAN link and disk utilization (see the previous description of factors to be aware of when load balancing is applied). Also, any change in the IP address of a client (which can happen when working in DHCP environments) may cause the client to switch to another branch WAE, which can cause the client to experience reduced performance until the client's working set is retrieved into the new cache.

Hashing that is based on a client IP address does not guarantee any locality of the hash key. For example, clients from the same subnet (which are likely to share and collaborate on the same content) may be assigned two different hash numbers and may be redirected to different branch WAEs, while clients from different subnets may be assigned the same hash number and may be redirected to the same branch WAE. Hashing that is based on a client IP address does guarantee consistency. For example, a client using the same IP address is redirected to the same branch WAE.

In the service farm, a lead WAE is chosen to build the hash table that distributes the load between the available WAEs. The lead WAE distributes the buckets evenly. The source IP address is hashed and the resulting bucket determines the WAE that will handle the packet.

WCCP supports redirection by mask value assignments. This method relies on masking to make redirection decisions. The decisions are made using special hardware support in the WCCP-enabled router. This method can be very efficient because packets are switched by the hardware.

Note The masking method can only be used for load balancing with the Catalyst 3750, Catalyst 4500, and Catalyst 6500 series switches, Cisco 7600 series routers, and Cisco ASR 1000 series routers. And, the masking method can be used with the Cisco 2800, 3800, and 7200 series routers when they are running Cisco IOS release 12.4(20)T or later releases.

You must explicitly specify masking. You can specify two mask values based on the source or destination IP address of the packet. For WAAS, the default mask value is based on the source IP address. You can enable masks by using the default values or specifying a particular mask. The default mask values, specified in hexadecimal notation, are as follows:

- dst-ip-mask= 0x0
- src-ip-mask= 0xF00

You may specify the mask value with a maximum of seven bits. The WAE creates a table of the 2^7 (or 128) combinations, assigns the WAE IP addresses to them, and sends this table to the WCCP-enabled routers. The router uses this table to distribute the traffic among all the WAEs that are in the service group. Each packet that matches the WCCP service parameters is compared to this table and the packets are sent to the matching WAE.

In a service farm where the WAEs have different masks, the first WAE to establish two-way communication with the routers determines the farm's mask. All other WAEs cannot join the farm unless they are configured with the same mask.

Masking is typically used at the data center, where you can take advantage of the hardware accelerated WCCP redirection capabilities of switches such as the Catalyst 6500 series switches. At the data center, the load balancing goal should be to have all connections originating from a given client subnet (typically equivalent to a branch) go to one data center WAE, to improve data redundancy elimination

(DRE) compression performance. Also, mask assignment on the Catalyst 6500 series switches uses the ACL TCAM. When combined with WCCP redirect lists, mask assignment can use a large portion of the TCAM. To minimize TCAM usage, use a mask with fewer care bits.

Given these considerations, beginning with WAAS version 4.2.1, the default mask has been changed from src-ip-mask 0x1741 and dst-ip-mask 0x0 (in 4.1x versions) to src-ip-mask 0xF00 and dst-ip-mask 0x0 (in 4.2.1 and later versions). The current source IP mask uses only 4 care bits rather than the 6 care bits used by the old mask.

With a typical data center WCCP interception configuration (ingress interception with service 61 on the WAN, ingress interception with service 62 on the LAN), this mask load balances /24 branch subnets (it extracts the last 4 bits of /24 subnets). Connections from one branch subnet will be pinned to one data center WAE. If your network has a different distribution of IP addresses (for example, /16 subnets), you should configure a mask that extracts bits from the /16 network part of the address, for example, src-ip-mask 0xF0000. Similarly, if some branches generate more traffic than others, you may want to create a mask that also extracts bits from the host part of the address, for example, 0xF03.

Information About Packet-Forwarding Methods

A WCCP-enabled router redirects intercepted TCP segments to a WAE using one of the following two packet-forwarding methods:

- Generic routing encapsulation (GRE)—Allows packets to reach the WAE even if there are any number of routers in the path to the WAE.
- Layer 2 redirection—Allows packets to be switched at Layer 2 (MAC layer) and reach the WAE.

Table 5-2 describes the packet-forwarding methods.

Table 5-2 Packet-Forwarding Methods

Packet-Forwarding Method	Load-Balancing Method: Hashing	Load-Balancing Method: Masking
GRE (Layer 3)	Packet redirection is completely handled by the router software.	Packet redirection is handled by the router software. We do not recommend using mask assignment when GRE is being used as the packet-forwarding method.
Layer 2 redirection	First redirected packet is handled by the router software; all subsequent redirected packets are handled by the router hardware.	All packets are handled by the router hardware (currently supported only on the Catalyst 6500 series switches or Cisco 7600 series routers because special hardware is required).

The redirection mode is controlled by the branch WAE. The first branch WAE that joins the WCCP service group decides the forwarding method (GRE or Layer 2 redirection) and the assignment method (hashing or masking). The term *mask assignment* refers to WCCP Layer 2 Policy Feature Card 2 (PFC2) input redirection.

If masking is selected with WCCP output redirection, then the branch WAE falls back to the original hardware acceleration that is used with the Multilayer Switch Feature Card (MSFC) and the Policy Feature Card (PFC).

For example, WCCP filters the packets to determine which redirected packets have been returned from the branch WAE and which ones have not. WCCP does not redirect the ones that have been returned because the branch WAE has determined that the packets should not be processed. WCCP Version 2 returns packets that the branch WAE does not service to the same router from which they were transmitted.

This section contains the following topics:

Reasons for Packet Rejection and Return

A branch WAE rejects packets and initiates packet return for the following reasons:

- The WAE is filtering out certain conditions that make processing packets unproductive, for example, when IP authentication has been turned on.
- You have configured a static bypass list or interception ACL on the branch WAE.

Note The packets are redirected to the source of the connection between the WCCP-enabled router and the branch WAE. Depending on the Cisco IOS software version used, this source could be either the address of the outgoing interface or the router IP address. In the latter case, it is important that the branch WAE has the IP address of the WCCP-enabled router stored in the router list. For more information on router lists, see the "Configuring and Viewing WCCP Router Lists for WAEs" section on page 5-26.

Cisco Express Forwarding (CEF) is required for WCCP and must be enabled on the router.

WCCP also allows you to configure multiple routers in a router list to support a particular WCCP service (for example, CIFS redirection).

Layer 3 GRE as a Packet-Forwarding Method

A WCCP-enabled router redirects intercepted requests to a WAE and can encapsulate the packets using GRE. This method for forwarding packets allows packets to reach the WAE even if there are routers in the path to the WAE. Packet redirection is handled entirely by the router software.

GRE allows datagrams to be encapsulated into IP packets at the WCCP-enabled router and then redirected to a WAE (the transparent proxy server). At this intermediate destination, the datagrams are decapsulated and then handled by the WAAS software. If the request cannot be handled locally, the origin server may be contacted by the associated WAE to complete the request. In doing so, the trip to the origin server appears to the inner datagrams as one hop. The redirected traffic using GRE usually is referred to as GRE tunnel traffic. With GRE, all redirection is handled by the router software.

With WCCP redirection, a Cisco router does not forward the TCP SYN packet to the destination because the router has WCCP enabled on the destination port of the connection. Instead, the WCCP-enabled router encapsulates the packet using GRE tunneling and sends it to the WAE that has been configured to accept redirected packets from this WCCP-enabled router.

After receiving the redirected packet, the WAE does the following:

1. Strips the GRE layer from the packet.
2. Decides whether it should accept this redirected packet and process the request for content or deny the redirected packet as follows:

a. If the WAE decides to accept the request, it sends a TCP SYN ACK packet to the client. In this response packet, the WAE uses the IP address of the original destination (origin server) that was specified as the source address so that the WAE can be invisible (transparent) to the client; it pretends to be the destination that the TCP SYN packet from the client was trying to reach.

b. If the WAE decides not to accept the request, it reencapsulates the TCP SYN packet in GRE, and sends it back to the WCCP-enabled router. The router understands that the WAE is not interested in this connection and forwards the packet to its original destination (that is, the origin server).

Layer 2 Redirection as a Packet-Forwarding Method

Layer 2 redirection is accomplished when a WCCP-enabled router or switch takes advantage of internal switching hardware that either partially or fully implements the WCCP traffic interception and redirection functions at Layer 2. This type of redirection is currently supported only with the Catalyst 6500 series switches and Cisco 7200 and 7600 series routers. With Layer 2 redirection, the first redirected traffic packet is handled by the router software. The rest of the traffic is handled by the router hardware. The branch WAE instructs the router or switch to apply a bit mask to certain packet fields, which in turn provides a mask result or index mapped to the branch WAE in the service group in the form of a mask index address table. The redirection process is accelerated in the switching hardware, making Layer 2 redirection more efficient than Layer 3 GRE.

Note WCCP is licensed only on the WAE and not on the redirecting router. WCCP does not interfere with normal router or switch operations.

Information About WCCP Flow Redirection on WAEs

Flow protection reduces the impact on existing client TCP connections when branch WAEs are added and removed from a service group. By default, WCCP flow redirection is disabled on a WAE. The client impact is reduced because of flow protection in the following situations, typical in large WCCP service farms:

- WAAS network expansion—When branch WAEs are added to the service group, the newly started branch WAEs receives traffic that was previously processed by a different branch WAE. It forwards the traffic to the relevant branch WAE for continued processing. New connections are processed by the new branch WAE.
- Branch WAE replacement following a failure—When a branch WAE fails, another branch WAE may receive traffic that was previously processed by either that branch WAE or the origin file server. The receiving branch WAE operates according to the previous two use cases.

Without flow protection, established client connections are broken through a TCP RESET in the situations listed earlier. Flow protection applies to all supported WCCP services and cannot be configured on a per-service basis.

To enable flow protection for a specified time period, use the **wccp flow-redirect enable timeout** *seconds* global configuration command. After the timeout period, flow protection ceases. If you do not specify the timeout option, flow protection is enabled indefinitely.

Note Network designs that require redirected frames to be returned to the originating router are not compatible with the WCCP flow protection feature.

Configuring or Viewing the WCCP Settings on WAEs

This section describes how to configure or view WCCP settings on WAEs that are configured as application accelerators and are not part of an AppNav Cluster (WAEs that are part of an AppNav Cluster use only the appnav-controller interception method). If you want to configure or view the WCCP settings on WAEs configured as AppNav Controllers, see "Configuring or Viewing the WCCP Settings on ANCs" section on page 5-22.

Device group configuration is not possible beginning with WAAS version 5.0. However, you can use the **Copy Settings** taskbar icon in the configuration window to copy the settings to other devices in your network. To ensure consistency, we recommend that you copy the same WCCP settings to all devices in the same WCCP service farm.

Note Before you do the procedure in this section, you should have already completed a basic WCCP configuration for your WAAS network that includes the configuration of the TCP promiscuous mode service as described in the *Cisco Wide Area Application Services Quick Configuration Guide*.

To modify the WCCP settings for a WAE, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Interception > Interception Configuration**. The Interception Configuration window appears. (See Figure 5-3.)

Note If you are configuring a device using a WAAS version earlier than 5.0, choose **Configure > Interception > WCCP > Settings** to configure WCCP settings. The configuration window looks different but has similar settings.

Figure 5-3 Interception Configuration Window for WAE

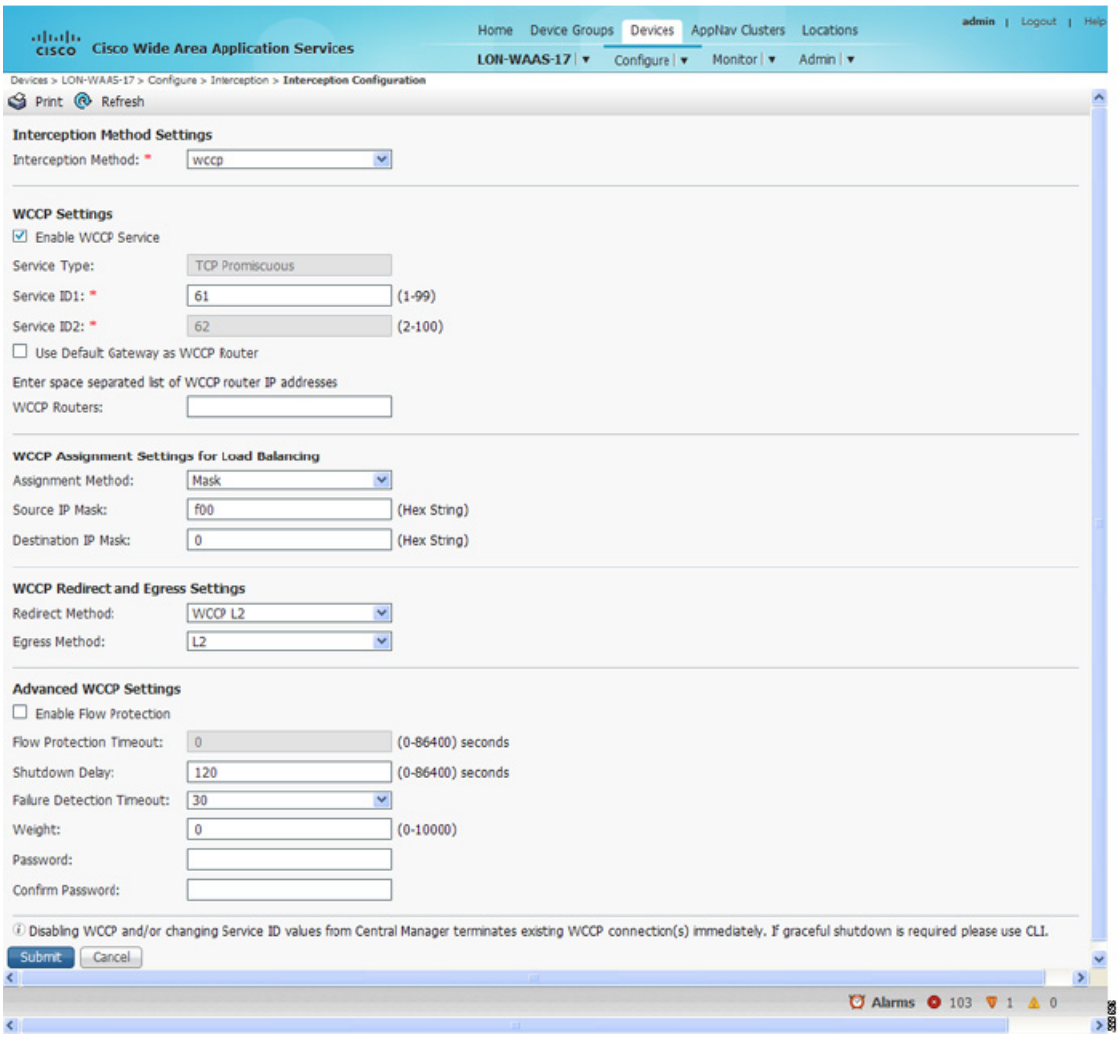

Step 3 Check the current settings for the chosen device:

- To keep the current settings and to close the window, click **Cancel**.
- To remove the current settings, click the **Remove Settings** taskbar icon.
- To modify the current settings, change the current setting as described in the rest of this procedure.
- To copy the settings to other WAEs in your network, click the **Copy Settings** taskbar icon. The Copy Interception Settings window opens where you can select other WAEs to which the interception settings can be copied. You can copy all settings or you can exclude the router list and enable the WCCP service. Click **OK** to copy the settings to the selected WAEs devices.

By default, WCCP is disabled on a WAE. However, as part of the initial configuration of WCCP in your WAAS network, you should have enabled WCCP Version 2 on your WAEs (the branch WAE and the data center WAE) as well as on the routers in the data center and branch office that will be transparently redirecting requests to these WAEs. For information about how to perform a basic WCCP configuration in your WAAS network, see the *Cisco Wide Area Application Services Quick Configuration Guide*.

Step 4 From the Interception Method drop-down list, choose **wccp** to enable the WCCP interception method. If you change this setting from any setting other than None, you must click the **Submit** button to update the window with the proper fields for configuring WCCP. (The Interception Method drop-down list is not shown for devices using WAAS versions earlier than 5.0.)

Step 5 Check the **Enable WCCP Service** check box to enable WCCP Version 2 on the chosen device, or uncheck the check box to disable WCCP on the chosen device.

Note Ensure that the routers used in the WCCP environment are running a version of the Cisco IOS software that also supports the WCCP Version 2.

Note If you use the Central Manager to disable WCCP on a WAAS device, the Central Manager immediately shuts down WCCP and closes any existing connections, ignoring the setting configured by the **wccp shutdown max-wait** global configuration command. If you want to gracefully shut down WCCP connections, use the **no enable** WCCP configuration command on the WAAS device.

Step 6 In the Service ID1 field, specify the first service ID of the WCCP service pair. After you submit, the Service ID2 field is filled in with the second service ID of the pair, which is one greater than Service ID1. For WAEs with version 4.4.1 or later, you can change the WCCP service IDs from the default of 61/62 to a different pair of numbers, which allows a router to support multiple WCCP farms because the WAEs in different farms can use different service IDs. (The Service ID fields are not shown for devices using WAAS versions earlier than 4.4 and the service IDs are fixed at 61/21.)

The router service priority varies inversely with the service ID. The service priority of the default service IDs 61/62 is 34. If you specify a lower service ID, the service priority is higher than 34; if you specify a higher service ID, the service priority is lower than 34.

Step 7 Check the **Use Default Gateway as WCCP Router** check box to use the default gateway of the WAE device as the router to associate with the WCCP TCP promiscuous mode service. Alternatively, you can uncheck this box and specify a list of one more routers by their IP addresses, separated by spaces. The Central Manager assigns the router list number, which is displayed next to the router list field after the page is submitted. As part of the initial configuration of your WAAS network, you may have already created a WCCP router list with the setup utility, as described in the *Cisco Wide Area Application Services Quick Configuration Guide*. For more information about WCCP router lists, see the "Configuring and Viewing WCCP Router Lists for WAEs" section on page 5-26.

Note Checking or unchecking this check box, changing the router list, or submitting the WCCP page removes any other existing router lists that are not assigned to the WCCP service, including router lists configured by the setup utility or through the CLI.

Step 8 (Optional) To force WCCP to use the configured assignment method only, check the **Only Use Selected Assignment Method** check box. You can specify only one load-balancing method (hashing or masking) per WCCP service in a branch WAE service group. (This check box is shown only for devices using WAAS versions earlier than 4.4.)

Note If you check the Only Use Selected Assignment Method check box, the WAE only joins a WCCP farm if the assignment method configured on the WAE is supported by the router. If you do not check the Only Use Selected Assignment Method check box, the WAE uses the assignment method that the router supports, even if the WAE is configured differently from the router.

Step 9 (Optional) From the Assignment Method drop-down list, choose the type of WAE load-balancing assignment method to use (for more information, see the "Information About Load Balancing and WAEs" section on page 5-11):

- Choose **Hash** to use the hash method (the default for devices using WAAS versions earlier than 5.0). Follow Steps 10 and 11 to define how the hash works, and skip to Step 13 because the mask settings are not used.
- Choose **Mask** to use the mask method (the default for devices using WAAS versions 5.0 or later). Skip to Step 12 to define the service mask.

Step 10 (Optional) To define the load-balancing hash for WCCP service ID1 on the source IP address, check the **Hash on Source IP** check box. This check box is shown only if the hash assignment method is used.

Step 11 (Optional) To define the load-balancing hash for WCCP service ID1 on the destination IP address, check the **Hash on Destination IP** check box. This check box is shown only if the hash assignment method is used.

Step 12 (Optional) To use a custom service mask, enter different mask values in the WCCP Assignment Settings for Load Balancing area, overwriting the default mask settings. If you do not change these settings, the defaults are used. Define the custom mask as follows:

- In the Source IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, FE000000) used to match the packet source IP address. The range is 00000000–FE000000. The default is F00.
- In the Destination IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, FE000000) used to match the packet destination IP address. The range is 0000000–FE000000. The default is 0.

Note If you apply the default mask to a WAE running version 4.1.x or earlier, the mask is different from the default mask (0x1741) set under software version 4.1.x and earlier.

If the WAE detects that its configured mask is not the same as advertised by one or more routers in the farm, it is not allowed to join the farm and a major alarm is raised ("Configured mask mismatch for WCCP"). This alarm can occur when a WAE is trying to join a farm that already has other WAEs and these other WAEs are configured with a different mask. The routers do not allow other WAEs to join the farm unless they advertise the same mask. To correct this alarm, ensure that all WAEs in the farm are configured with the same mask. This alarm is cleared when the WAE's configured mask matches the mask of all the routers in the farm.

Step 13 From the Redirect Method drop-down list, choose the type of packet redirection (forwarding) method to use:

- **WCCP GRE** (the default for devices using WAAS versions earlier than 5.0) to use Layer 3 GRE packet redirection.
- **WCCP L2** (the default for devices using WAAS versions 5.0 or later) to permit the WAE to receive transparently redirected traffic from a WCCP Version 2-enabled switch or router if the WAE has a Layer 2 connection with the device and the device is configured for Layer 2 redirection. For more information, see the "Information About Packet-Forwarding Methods" section on page 5-14.

Step 14 From the Return Method drop-down list, choose the type of method to use to return nonoptimized (bypassed) packets to the router:

- **WCCP GRE** (the default) to use GRE packet return.
- **WCCP L2** to use Layer 2 rewriting for packet return.

(The Return Method drop-down list is shown only for devices using WAAS versions earlier than 5.0. For later WAAS versions, the return method is set the same as the redirect method.)

Step 15 (Optional) From the Egress Method drop-down list, choose the method to use to return optimized packets to the router or switch:

- **Generic GRE** (available and set as the default only if Redirect Method is WCCP GRE)
- **IP Forwarding**
- **L2** (available and set as the default only if Redirect Method is WCCP L2)
- **WCCP GRE** (available only if Redirect Method is WCCP GRE)

For devices using WAAS versions earlier than 5.0, the choices are as follows: IP Forwarding (the default), WCCP Negotiated Return, or Generic GRE. For more details on choosing the egress method, see the "Configuring Egress Methods for WCCP Intercepted Connections" section on page 5-29.

Step 16 (Optional) Modify the current advanced settings in the Advanced WCCP Settings area as follows:

a. Check the **Enable Flow Protection** check box to keep the TCP flow intact and to avoid overwhelming the device when it comes up or is reassigned new traffic. For more information, see the "Information About WCCP Flow Redirection on WAEs" section on page 5-16. Flow protection is disabled by default.

b. In the Flow Protection Timeout field, specify the amount of time (in seconds) that flow protection should be enabled. The default is 0, which means it stays enabled with no timeout. (The Flow Protection Timeout field is not shown for devices using WAAS versions earlier than 5.0.)

c. In the Shutdown Delay field, specify the maximum amount of time (in seconds) that the chosen device waits to perform a clean shutdown of WCCP. The default is 120 seconds.

The WAE does not reboot until either all connections have been serviced or the maximum wait time (specified through this Shutdown Delay field) has elapsed for WCCP.

d. In the Failure Detection Timeout drop-down list, choose the failure detection timeout value (9, 15, or 30 seconds). The default is 30 seconds and is the only value supported on WAAS versions prior to 4.4.1. This failure detection value determines how long it takes the router to detect a WAE failure. (The Failure Detection Timeout field is not shown for devices using WAAS versions earlier than 4.4.)

The failure detection timeout value is negotiated with the router and takes effect only if the router also has the variable timeout capability. If the router has a fixed timeout of 30 seconds and you have configured a failure detection value on the WAE other than the default 30 seconds, the WAE is not able to join the farm and an alarm is raised ("Router unusable" with a reason of "Timer interval mismatch with router").

e. In the Weight field, specify the weight value that is used for load balancing. The weight value ranges from 0 to 10000. If the total of all the weight values of the WAEs in a service group is less than or equal to 100, then the weight value represents a literal percentage of the total load redirected to the device for load-balancing purposes. For example, a WAE with a weight of 10 receives 10 percent of the total load in a service group where the total of all weight values is 50. If a WAE in such a service group fails, the other WAEs still receive the same load percentages as before the failure; they will not receive the load allocated to the failed WAE.

If the total of all the weight values of the WAEs in a service group is between 101 and 10000, then the weight value is treated as a fraction of the total weight of all the active WAEs in the service group. For example, a WAE with a weight of 200 receives 25 percent of the total load in a service group where the total of all the weight values is 800. If a WAE in such a service group fails, the other WAEs will receive the load previously allocated to the failed WAE. The failover handling is different than if the total weights are less than or equal to 100.

By default, weights are not assigned and the traffic load is distributed evenly between the WAEs in a service group.

f. In the Password field, specify the password to be used for secure traffic between the WAEs within a cluster and the router for a specified service. Be sure to enable all other WAEs and routers within the cluster with the same password. Passwords must not exceed eight characters in length. Do not use the following characters: space, backwards single quote (`), double quote ("), pipe (|), or question mark (?). Reenter the password in the Confirm Password field.

Note For information about how to use the CLI to specify the service group password on a router, see the "Setting a Service Group Password on a Router" section on page 5-10.

Step 17 Click **Submit** to save the settings.

To configure WCCP settings from the CLI, you must first set the interception method to WCCP by using the **interception-method** global configuration command, and then you can use the **wccp flow-redirect**, **wccp router-list**, **wccp shutdown**, and **wccp tcp-promiscuous** global configuration commands.

For more information about a graceful shut down of WCCP Version 2 on WAEs, see the "Configuring WAEs for a Graceful Shutdown of WCCP" section on page 5-26.

Configuring or Viewing the WCCP Settings on ANCs

This section describes how to configure or view WCCP settings on WAAS devices configured as AppNav Controllers (ANCs). Typically, you configure ANCs and their settings through the AppNav Clusters window in the Central Manager, which includes WCCP settings, so you do not need to configure the WCCP settings outside the AppNav Cluster context as described in this section.

If you want to configure or view the WCCP settings on WAEs configured as application accelerators, see the "Configuring or Viewing the WCCP Settings on WAEs" section on page 5-17. To configure interception settings on WAEs operating as WAAS nodes for an AppNav Controller, see the "Configuring AppNav Interception" section on page 5-56.

Device group configuration is not possible beginning with WAAS version 5.0. However, you can use the **Copy Settings** taskbar icon in the configuration window to copy the settings to other devices in your network. To ensure consistency, we recommend that you copy the same WCCP settings to all devices in the same WCCP service farm.

To modify the WCCP settings for an ANC, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Interception > Interception Configuration**. The Interception Configuration window appears. (See Figure 5-3.)

Figure 5-4 Interception Configuration Window for ANC

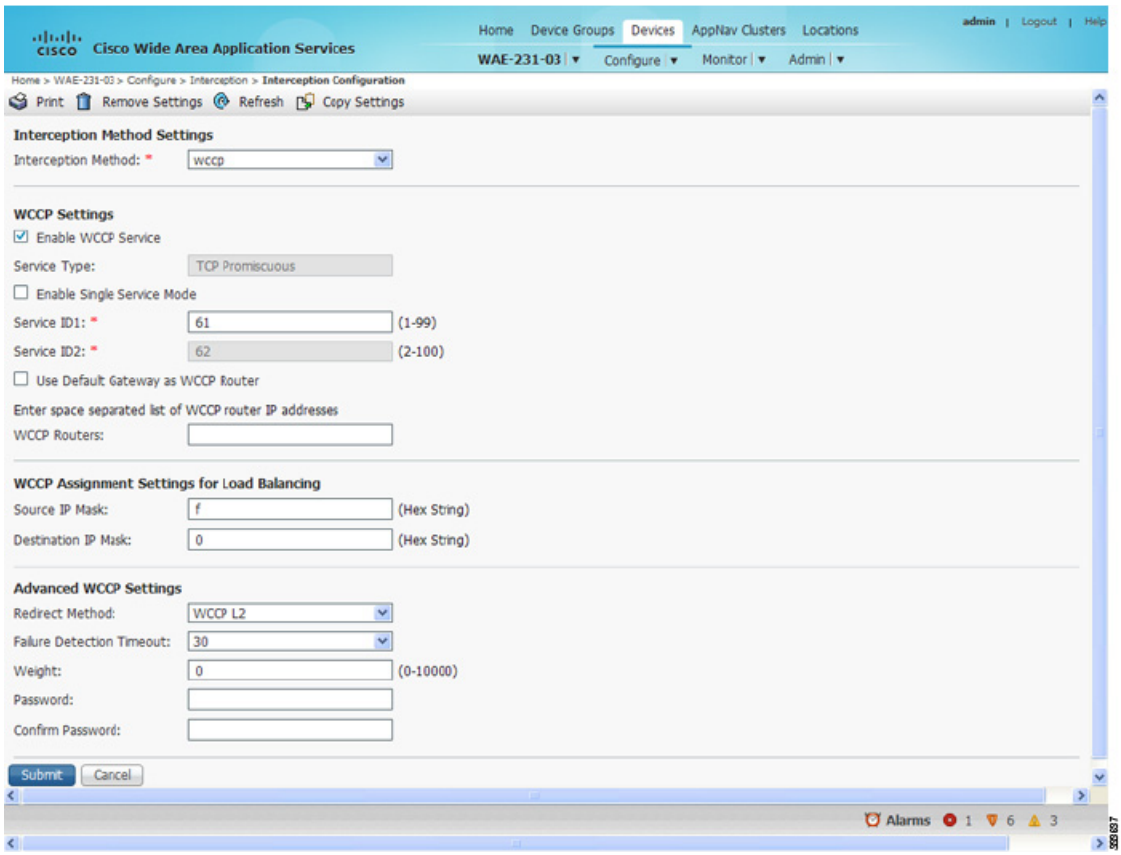

Step 3 Check the current settings for the chosen device:

- To keep the current settings and to close the window, click **Cancel**.
- To remove the current settings, click the **Remove Settings** taskbar icon.
- To modify the current settings, change the current setting as described in the rest of this procedure.
- To copy the settings to other WAEs in your network, click the **Copy Settings** taskbar icon. The Copy Interception Settings window opens where you can select other WAEs to which the interception settings can be copied. You can copy all settings or you can exclude the router list and enable the WCCP service. Click **OK** to copy the settings to the selected WAEs devices.

By default, WCCP is disabled on a WAE. However, as part of the initial configuration of WCCP in your WAAS network, you should have enabled WCCP Version 2 on your WAEs (the branch WAE and the data center WAE) as well as on the routers in the data center and branch office that will be transparently redirecting requests to these WAEs. For information about how to perform a basic WCCP configuration in your WAAS network, see the *Cisco Wide Area Application Services Quick Configuration Guide*.

Step 4 From the Interception Method drop-down list, choose **wccp** to enable the WCCP interception method. If you change this setting from any setting other than None, you must click the **Submit** button to update the window with the proper fields for configuring WCCP.

Step 5 Check the **Enable WCCP Service** check box to enable WCCP Version 2 on the chosen device, or uncheck the check box to disable WCCP on the chosen device.

Note Ensure that the routers used in the WCCP environment are running a version of the Cisco IOS software that also supports the WCCP Version 2.

Note If you use the Central Manager to disable WCCP on a WAAS device, the Central Manager immediately shuts down WCCP and closes any existing connections, ignoring the setting configured by the **wccp shutdown max-wait** global configuration command. If you want to gracefully shut down WCCP connections, use the **no enable** WCCP configuration command on the WAAS device.

Step 6 (Optional) You can enable single service mode by checking the **Enable Single Service Mode** check box (the default). Single service mode simplifies configuration by using the same service ID for incoming and outgoing traffic, which is possible only with an AppNav deployment because it can handle asymmetric traffic flows.

Step 7 In the Service ID1 field, specify the service ID of the WCCP service.

If the Enable Single Service Mode check box is unchecked, a pair of WCCP service IDs are required and the Service ID2 field is filled in with the second service ID of the pair, which is one greater than Service ID1. The default service IDs are 61 and 62. You can change the WCCP service IDs from the default of 61/62 to a different pair of numbers, which allows a router to support multiple WCCP farms because the ANCs in different farms can use different service IDs.

The router service priority varies inversely with the service ID. The service priority of the default service IDs 61/62 is 34. If you specify a lower service ID, the service priority is higher than 34; if you specify a higher service ID, the service priority is lower than 34.

Step 8 Check the **Use Default Gateway as WCCP Router** check box to use the default gateway of the WAE device as the router to associate with the WCCP TCP promiscuous mode service. Alternatively, you can uncheck this box and specify a list of one more routers by their IP addresses, separated by spaces. The Central Manager assigns the router list number, which is displayed next to the router list field after the page is submitted. As part of the initial configuration of your WAAS network, you may have already created a WCCP router list with the setup utility, as described in the *Cisco Wide Area Application Services Quick Configuration Guide*. For more information about WCCP router lists, see the "Configuring and Viewing WCCP Router Lists for WAEs" section on page 5-26.

Note Checking or unchecking this check box, changing the router list, or submitting the WCCP page removes any other existing router lists that are not assigned to the WCCP service, including router lists configured by the setup utility or through the CLI.

Step 9 (Optional) To use a custom service mask, enter different mask values in the WCCP Assignment Settings for Load Balancing area, overwriting the default mask settings. If you do not change these settings, the defaults are used. Define the custom mask as follows (for more information, see the "Information About Load Balancing and WAEs" section on page 5-11):

- In the Source IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, FE000000) used to match the packet source IP address. The range is 00000000–FE000000. The default is F.
- In the Destination IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, FE000000) used to match the packet destination IP address. The range is 0000000–FE000000. The default is 0.

If the WAE detects that its configured mask is not the same as advertised by one or more routers in the farm, it is not allowed to join the farm and a major alarm is raised ("Configured mask mismatch for WCCP"). This alarm can occur when a WAE is trying to join a farm that already has other WAEs and these other WAEs are configured with a different mask. The routers do not allow other WAEs to join the farm unless they advertise the same mask. To correct this alarm, ensure that all WAEs in the farm are configured with the same mask. This alarm is cleared when the WAE's configured mask matches the mask of all the routers in the farm.

Step 10 (Optional) Modify the current advanced settings in the Advanced WCCP Settings area as follows:

a. From the Redirect Method drop-down list, choose the type of packet redirection (forwarding) method to use:

- **WCCP GRE** to use Layer 3 GRE packet redirection.
- **WCCP L2** (the default) to permit the WAE to receive transparently redirected traffic from a WCCP Version 2-enabled switch or router if the WAE has a Layer 2 connection with the device and the device is configured for Layer 2 redirection. For more information, see the "Information About Packet-Forwarding Methods" section on page 5-14.

The return method is the same as the redirect method. The egress method is generic GRE when the WCCP GRE redirect method is chosen or WCCP L2 return when the WCCP L2 redirect method is chosen.

b. In the Failure Detection Timeout drop-down list, choose the failure detection timeout value (3, 6, 9, 15, or 30 seconds). The default is 30 seconds and is the only value supported on WAAS versions prior to 4.4.1. This failure detection value determines how long it takes the router to detect a WAE failure.

The failure detection timeout value is negotiated with the router and takes effect only if the router also has the variable timeout capability. If the router has a fixed timeout of 30 seconds and you have configured a failure detection value on the WAE other than the default 30 seconds, the WAE is not able to join the farm and an alarm is raised ("Router unusable" with a reason of "Timer interval mismatch with router").

c. In the Weight field, specify the weight value that is used for load balancing. The weight value ranges from 0 to 10000. If the total of all the weight values of the WAEs in a service group is less than or equal to 100, then the weight value represents a literal percentage of the total load redirected to the device for load-balancing purposes. For example, a WAE with a weight of 10 receives 10 percent of the total load in a service group where the total of all weight values is 50. If a WAE in such a service group fails, the other WAEs still receive the same load percentages as before the failure; they will not receive the load allocated to the failed WAE.

If the total of all the weight values of the WAEs in a service group is between 101 and 10000, then the weight value is treated as a fraction of the total weight of all the active WAEs in the service group. For example, a WAE with a weight of 200 receives 25 percent of the total load in a service group where the total of all the weight values is 800. If a WAE in such a service group fails, the other WAEs will receive the load previously allocated to the failed WAE. The failover handling is different than if the total weights are less than or equal to 100.

By default, weights are not assigned and the traffic load is distributed evenly between the WAEs in a service group.

d. In the Password field, specify the password to be used for secure traffic between the WAEs within a cluster and the router for a specified service. Be sure to enable all other WAEs and routers within the cluster with the same password. Passwords must not exceed eight characters in length. Do not use the following characters: space, backwards single quote (`), double quote ("), pipe (|), or question mark (?). Reenter the password in the Confirm Password field.

Note For information about how to use the CLI to specify the service group password on a router, see the "Setting a Service Group Password on a Router" section on page 5-10.

Step 11 Click **Submit** to save the settings.

To configure WCCP settings from the CLI, you must first set the interception method to WCCP by using the **interception-method** global configuration command, and then you can use the **wccp router-list** and **wccp tcp-promiscuous** global configuration commands.

Configuring and Viewing WCCP Router Lists for WAEs

You can configure and view one router list from the Central Manager through the WCCP settings (see the "Configuring or Viewing the WCCP Settings on WAEs" section on page 5-17). The Central Manager supports only a single router list assigned to the WCCP service and removes any other existing router lists that may be configured through the CLI if you use the Central Manager to configure a router list, check or uncheck the Use Default Gateway check box in the WCCP settings page, or submit the WCCP settings page. If you want to configure a router list through the CLI, you can use the **wccp router-list** global configuration command.

Note WCCP must be enabled before you can use the WCCP global configuration commands.

To delete a router list, use the **no wccp router-list** global configuration command.

To view an unassigned router list configured by the **wccp router-list** command, use the **show running-config wccp** EXEC command.

Configuring WAEs for a Graceful Shutdown of WCCP

To prevent broken TCP connections, the WAE performs a clean shutdown of WCCP after you disable WCCP Version 2 on a WAE or reload the WAE from the CLI. You can perform this task locally through the CLI on a device by entering the **no enable** WCCP configuration command.

The WAAS Central Manager also allows you to disable WCCP Version 2 on a WAE, but this does not perform a graceful shut down of WCCP connections. To disable WCCP immediately for a chosen device, uncheck the **Enable WCCP** check box in the WAAS Central Manager Interception Configuration window. (See Figure 5-3.)

Note If you use the Central Manager to disable WCCP on a WAAS device, the Central Manager immediately shuts down WCCP and closes any existing connections, ignoring the setting configured by the **wccp shutdown max-wait** global configuration command. If you want to gracefully shut down WCCP connections, use the **no enable** WCCP configuration command on the WAAS device.

During a graceful shut down, the WAE does not reboot until one of the following occurs:

- All the connections have been serviced.

- The maximum wait time (specified through the Shutdown Delay field in the WCCP Configuration Settings window or with the **wccp shutdown max-wait** command [by default, 120 seconds]) has elapsed for WCCP Version 2.

During a clean shutdown of WCCP, the WAE continues to service the flows that it is handling, but it starts to bypass new flows. When the number of flows goes down to zero, the WAE takes itself out of the group by having its buckets reassigned to other WAEs by the lead WAE. TCP connections can still be broken if the WAE crashes or is rebooted without WCCP being cleanly shut down.

You cannot shut down an individual WCCP service on a particular port on a WAE; you must shut down WCCP on the WAE. After WCCP is shut down on the WAE, the WAE preserves its WCCP configuration settings.

Configuring Static Bypass Lists for WAEs

Note Static bypass lists are supported only for devices (but not device groups) using WAAS versions earlier than 5.0 and are deprecated for such devices. Interception ACLS are recommended instead.

Using a static bypass allows traffic flows between a configurable set of clients and servers to bypass handling by the WAE. By configuring static bypass entries on the branch WAE, you can control traffic interception without modifying the router configuration. IP access lists may be configured separately on the router to bypass traffic without first redirecting it to the branch WAE. Typically, the WCCP accept list defines the group of servers that are accelerated (and the servers that are not). Static bypass can be used occasionally when you want to prevent WAAS from accelerating a connection from a specific client to a specific server (or from a specific client to all servers).

Note We recommend that you use ACLs on the WCCP-enabled router where possible, rather than using static bypass lists or interception ACLs on the WAEs, because that is the most efficient method to control traffic interception. If you decide to use static bypass lists or interception ACLs, we recommend using interception ACLs because they are more flexible and give better statistics about passed-through connections. For information about how to configure ACLs on a router, see the "Configuring IP Access Lists on a Router" section on page 5-9. For information about how to configure an interception ACL for a WAE, see the "Configuring Interception Access Control Lists" section on page 5-28.

To configure a static bypass list for a version 4.x WAE, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Interception > Bypass Lists**.

Step 3 In the taskbar, click the **Create New WCCP/Inline Bypass List** icon. The Creating new WCCP/Inline Bypass List window appears.

Step 4 In the Client Address field, enter the IP address for the client.

Step 5 In the Server Address field, enter the IP address for the server.

Step 6 Check **Submit** to save the settings.

To configure a static bypass list from the CLI, you can use the **bypass static** global configuration command.

Configuring Interception Access Control Lists

You can configure an interception ACL to control what incoming traffic across all interfaces is to be intercepted by an ANC or WAE device (on an ANC, the interception ACL is called an AppNav Controller interception ACL). Packets that are permitted by the ACL are intercepted by the device, and packets that are denied by the ACL are passed through without processing.

By configuring an interception ACL on the WAAS device, you can control traffic interception without modifying the router configuration. IP ACLs may be configured separately on the router to bypass traffic without first redirecting it to the WAAS device. Typically, the WCCP accept list defines the group of servers that are accelerated (and the servers that are not). Using an interception ACL allows you to easily bypass uninteresting traffic, for example in a pilot deployment where you do not want to modify the router configuration. Additionally, it allows you to more easily transition from a pilot to a production deployment by allowing and accelerating different kinds of traffic in phases.

An interception ACL can be used both with WCCP and inline interception.

When used with interface ACLs and WCCP ACLs, the interface ACL is applied first, the WCCP ACL is applied second, and then the interception ACL is applied last. Application policies defined on the WAE are applied after all ACLs have filtered the traffic.

An ANC that is also operating as a WAAS node can have both an AppNav Controller interception ACL to control what is intercepted by the ANC and an interception ACL to control what is accepted by the optimizing engine. A flow may be permitted by the AppNav Controller interception ACL and then subsequently rejected by the WAAS node interception ACL.

Note The interception ACL feature is mutually exclusive with static bypass lists. You cannot use both types of lists at the same time. We recommend that you use interception ACLs instead of static bypass lists. Static bypass lists are supported only for devices using WAAS versions earlier than 5.0.

To use an interception ACL, first define an ACL (see Chapter 9, "Creating and Managing IP Access Control Lists for WAAS Devices") and then apply it to a device. Interception ACLs are configured for individual devices only and not device groups.

To configure an interception ACL for an ANC or WAE device, follow these steps:

Step 1 Follow the instructions in Chapter 9, "Creating and Managing IP Access Control Lists for WAAS Devices" to create an ACL that you want to use for interception, but do not apply it to an interface.

Step 2 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 3 Choose **Configure > Interception > Interception Access List**.

Step 4 To configure a WAE interception ACL, click the arrow control next to the Interception Access List field to display a drop-down list of ACLs you have defined and choose an ACL to apply to WAE interception. Alternatively, you can enter an ACL name directly in the field and create it after you submit this page. If you type in this field, the drop-down list of displayed ACLs is filtered to show only entries beginning with entered text.

If you need to create or edit an ACL, click the **Go to IP ACL** link next to the field to take you to the IP ACL configuration window (this is the Configure > Network > TCP/IP Settings > IP ACL page).

Step 5 To configure an ANC interception ACL, click the arrow control next to the AppNav Controller Interception Access List field to display a drop-down list of ACLs you have defined and choose an ACL to apply to ANC interception. Alternatively, you can enter an ACL name directly in the field and create

it after you submit this page. If you type in this field, the drop-down list of displayed ACLs is filtered to show only entries beginning with entered text. This field is shown only on devices configured in appnav-controller mode.

If you need to create or edit an ACL, click the **Go to IP ACL** link to take you to the IP ACL configuration window (this is the Configure > Network > TCP/IP Settings > IP ACL page).

Step 6 Check **Submit** to save the settings.

Note In AppNav Controller interception ACLs, the **tcp ... established** extended ACL condition is not supported and is ignored if encountered.

To configure an interception ACL from the CLI, you can use the **ip access-list** and **interception access-list** global configuration commands. To configure an AppNav Controller interception ACL, use the **interception appnav-controller access-list** global configuration command.

You can determine if a connection was passed through by an interception ACL by using the **show statistics connection** EXEC command. Flows passed through by an interception ACL are identified with a connection type of "PT Interception ACL."

Additionally, the **show statistics pass-through** command "Interception ACL" counter reports the number of active and completed pass through flows due to an interception ACL.

You can use the **show ip access-list** command to view the individual ACL rules that are being matched.

Configuring Egress Methods for WCCP Intercepted Connections

This section contains the following topics:

Information About Egress Methods

The WAAS software supports the following egress methods for WCCP intercepted connections:

- IP forwarding
- WCCP GRE return (available only if the redirect method is WCCP GRE; called WCCP negotiated return for devices earlier than version 5.0)
- Generic GRE (available only if the redirect method is WCCP GRE)
- Layer 2 (available only if the redirect method is WCCP L2)

Note For ANCs the egress method is not configurable. The egress method that is used depends on the redirect method. The ANC uses generic GRE when the WCCP GRE redirect method is chosen, or Layer 2 when the WCCP L2 redirect method is chosen.

The default egress method is L2. This egress method sends optimized data out through a Layer 2 connection to the router. This method is available only if the redirect method is also set to WCCP L2, and is not available on devices using WAAS versions earlier than 5.0. The router must also support Layer 2 redirect. If you configure the WCCP GRE redirect method or switch between WCCP GRE and L2, the default egress method is set to IP Forwarding.

For devices with a WAAS version earlier than 5.0, the default egress method is IP forwarding. The IP forwarding egress method does not allow you to place WAEs on the same VLAN or subnet as the clients and servers, and it does not ensure that packets are returned to the intercepting router.

The WCCP GRE return and generic GRE egress methods allow you to place WAEs on the same VLAN or subnet as clients and servers. Repeating redirection is prevented by encapsulating the outgoing frames in the GRE frames. Cisco IOS routers handle these GRE frames as bypass frames and do not apply WCCP redirection. With the WCCP GRE return method, WAAS uses the router ID address as the destination for GRE frames; with the generic GRE method, WAAS uses the address of the router configured in the WAE router list.

This technique makes it possible to support redundant routers and router load balancing; WAAS makes a best effort to return frames back to the router from which they arrived, though this is not guaranteed. An exception is that if flow protection is enabled, the WAE is unable to return flow-protected traffic to the originating router because the router information is not available.

Note Network designs that require redirected frames to be returned to the originating router are not compatible with the WCCP flow-protection feature.

If you want to use this functionality with multiple routers connected to the WAAS network segment, you must ensure connectivity to the router ID address, for example, by configuring static routes. The router ID is the address of the first loopback interface or highest active physical interface. This address can be found in the output of the **show wccp routers** EXEC command.

WAAS applies the following logic in its router selection for WCCP GRE and generic GRE:

- When the WAAS software applies data redundancy elimination (DRE) and compression to a TCP flow, the number of packets that are sent out may be fewer. A single packet that carries optimized data may represent original data that was received in multiple packets redirected from different routers. That optimized data-carrying packet will egress from the WAE to the router that last redirected a packet to the WAE for that flow direction.
- When the WAE receives optimized data, the data may arrive in multiple packets from different routers. The WAAS software expands the optimized data back to the original data, which will be sent out as several packets. Those original data-carrying packets will egress from the WAE to the router that last redirected a packet to the WAE for that flow direction.

The WCCP GRE return and generic GRE egress methods are similar, but the generic GRE egress method is designed specifically to be used in deployments where the router or switch does hardware-accelerated processing of GRE packets, such as with the Cisco 7600 series router or the Catalyst 6500 series switch with the Supervisor Engine 32 or 720. Additionally, the generic GRE egress method returns packets to the intercepting router by using a GRE tunnel that you must configure on the router (the WAE end of the tunnel is configured automatically). The generic GRE egress method is supported only when the WCCP GRE interception method is used.

To use the generic GRE egress method, you must create an intercepting router list on the WAE (multicast addresses are not supported) and configure a GRE tunnel interface on each router. For details on configuring GRE tunnel interfaces on the routers, see the "Configuring a GRE Tunnel Interface on a Router" section on page 5-31.

Note For devices with WAAS versions earlier than 5.0, WCCP Version 2 is capable of negotiating the redirect method and the return method for intercepted connections. The WAAS software supports WCCP GRE and WCCP Layer 2 as WCCP-negotiated return methods. If WCCP negotiates a WCCP Layer 2 return, the WAE defaults to using IP forwarding as the egress method. The WAE also defaults to IP forwarding if the interception method is set to WCCP Layer 2 and you configure generic GRE as the egress method, which are not compatible. When the WAE defaults to IP forwarding, the WAE logs a minor alarm that is cleared when you correct the configuration so that the interception and egress methods are consistent. The output of the **show egress methods** EXEC command also displays a warning if the interception and egress methods are not consistent.

For devices with WAAS version 5.0, you must explicitly configure the egress method.

Configuring the Egress Method

To configure the egress method for WCCP-intercepted connections from the Central Manager, see the "Configuring or Viewing the WCCP Settings on WAEs" section on page 5-17.

To configure the egress method for WCCP GRE packet return from the CLI, use the **egress-method** WCCP configuration command:

```
WAE(config)# wccp tcp-promiscuous service-pair 61 62
WAE(config-wccp-service)# egress-method wccp-gre
```

To configure the egress method for L2 return from the CLI, use the **egress-method** WCCP configuration command:

```
WAE(config)# wccp tcp-promiscuous service-pair 61 62
WAE(config-wccp-service)# egress-method L2
```

To configure the generic GRE egress method from the CLI, configure an intercepting router list and configure the egress method, as follows:

```
WAE(config)# wccp router-list 1 192.168.68.98
WAE(config)# wccp tcp-promiscuous service-pair 61 62
WAE(config-wccp-service)# router-list-num 1
WAE(config-wccp-service)# egress-method generic-gre
```

The router list must contain the IP address of each intercepting router. Multicast addresses are not supported. Additionally, you must configure a GRE tunnel interface on each router. For details on configuring GRE tunnel interfaces on the routers, see the "Configuring a GRE Tunnel Interface on a Router" section on page 5-31.

To view the egress method that is configured and that is being used on a particular WAE, use the **show wccp egress** EXEC command. To view information about the egress method for each connection segment, use the **show statistics connection egress-methods** EXEC command.

To view the generic GRE tunnel statistics for each intercepting router, use the **show statistics generic-gre** EXEC command. To clear statistics information for the generic GRE egress method, use the **clear statistics generic-gre** EXEC command.

Configuring a GRE Tunnel Interface on a Router

If you plan to use the generic GRE egress method on the WAE, you must configure a GRE tunnel interface on each intercepting router. For ease of configuration, we recommend that you create a single multipoint tunnel on the router, instead of one point-to-point tunnel per WAE in the farm.

If you have only one WAE in the farm, you can use a point-to-point tunnel, however, ensure that the router is configured with no other tunnel that has the same tunnel source as the WAE tunnel.

Note On the Catalyst 6500 series switch with the Supervisor Engine 32 or 720, do not configure more than one GRE tunnel (multipoint or point-to-point) with the same tunnel source interface, otherwise, high switch CPU load can result.

The tunnel interface must have a Layer 3 source interface to which it is attached and this source interface must be the interface whose IP address is configured in the WAE's intercepting router list.

The tunnel interface must be excluded from WCCP interception to avoid routing loops when outbound interception is used. Use the **ip wccp redirect exclude in** command. You can always use this command because it does not cause any impact even when it is not needed, such as for inbound interception.

This section contains the following topics:

- Multipoint Tunnel Configuration, page 5-32
- Point-To-Point Tunnel Configuration, page 5-33

Multipoint Tunnel Configuration

Consider a deployment in which there are two intercepting routers and two WAEs in the farm. Each WAE configuration would look like the following example:

```
wccp router-list 1 192.168.1.1 192.168.2.1
wccp tcp-promiscuous service-pair 61 62
  router-list-num-1
  egress-method generic-gre
  redirect-method gre
  enable
```

Each router can configure a single multipoint GRE tunnel to the WAE farm.

The router 1 configuration would look like the following example:

```
interface gigabitEthernet 1/1
ip address 192.168.1.1 255.255.255.0
...
interface Tunnel1
ip address 12.12.12.1 255.255.255.0
tunnel source GigabitEthernet1/1
tunnel mode gre multipoint
ip wccp redirect exclude in
end
```

The router 2 configuration would look like the following:

```
interface Vlan815 1/0
ip address 192.168.2.1 255.255.255.0
...
interface Tunnel1
ip address 13.13.13.1 255.255.255.0
tunnel source vlan815
tunnel mode gre multipoint
ip wccp redirect exclude in
end
```

Note The tunnel interface is enabled for IP by provisioning an IP address, which allows it to process and forward transit packets. If you do not want to provision an IP address, the tunnel must be IP enabled by making it an IP unnumbered interface. This restricts the tunnel to be a point-to-point tunnel.

Point-To-Point Tunnel Configuration

This section describes how to configure a point-to-point tunnel for a single WAE instead of a multipoint tunnel on the router. A point-to-point tunnel is enabled for IP either by making it unnumbered or by giving it an IP address. The unnumbered method is shown in the following example router configuration:

```
interface gigabitEthernet 1/1
ip address 192.168.1.1 255.255.255.0
...
! Tunnel1 is an unnumbered point-to-point tunnel towards WAE1
interface Tunnel1
ip unnumbered GigabitEthernet1/1
tunnel source GigabitEthernet1/1
! tunnel destination is the IP address of WAE1
tunnel destination 10.10.10.10
ip wccp redirect exclude in
end
```

Using Policy-Based Routing Interception

This section contains the following topics:

Information About Policy-Based Routing

Policy-based routing (PBR), introduced in Cisco IOS Release 11.0, allows you to implement policies that selectively cause packets to take specific paths in the network.

PBR also provides a method to mark packets so that certain kinds of traffic receive differentiated, preferential service when used in combination with queuing techniques enabled through the Cisco IOS software. These queuing techniques provide an extremely powerful, simple, and flexible tool to network managers who implement routing policies in their networks.

PBR enables the router to put packets through a route map before routing them. When configuring PBR, you must create a route map that specifies the match criteria and the resulting action if all of the match clauses are met. You must enable PBR for that route map on a particular interface. All packets arriving on the specified interface matching the match clauses will be subject to PBR.

One interface can have only one route map tag; but you can have several route map entries, each with its own sequence number. Entries are evaluated in order of their sequence numbers until the first match occurs. If no match occurs, packets are routed as usual.

```
Router(config-if)# ip policy route--tag
```

The route map determines which packets are routed next.

You can enable PBR to establish a route that goes through WAAS for some or all packets. WAAS proxy applications receive PBR-redirected traffic in the same manner as WCCP redirected traffic, as follows:

1. In the branch office, define traffic of interest on the branch office router (Edge-Router1) as follows:
   a. Specify which traffic is of interest to the LAN interface (ingress interface) on Edge-Router1.

      Use extended IP access lists to define traffic of interest (traffic from all or filtered local source addresses to any or filtered destination address).
   b. Specify which traffic is of interest to the WAN interface (egress interface) on Edge-Router1.

      Use extended IP access lists to define traffic of interest (traffic from all or filtered local source addresses from any or filtered remote addresses).
2. In the data center, specify which traffic is of interest to the data center router (Core-Router1) as follows:
   a. Specify which traffic is of interest to the LAN interface (ingress interface) on Core-Router1.

      Use extended IP access lists to define traffic of interest (traffic from all or filtered local source addresses to any or filtered destination address).
   b. Specify which traffic is of interest to the WAN interface (egress interface) on Core-Router1.

      Use extended IP access lists to define traffic of interest (traffic from all or filtered local source addresses from any or filtered remote addresses).
3. In the branch office, create route maps on Edge-Router1, as follows:
   a. Create a PBR route map on the LAN interface of Edge-Router1.
   b. Create a PBR route map on the WAN interface of Edge-Router1.
4. In the data center, create route maps on Core-Router1, as follows:
   a. Create a PBR route map on the LAN interface of Core-Router1.
   b. Create a PBR route map on the WAN interface of Core-Router1.
5. In the branch office, apply the PBR route maps to Edge-Router1.
6. In the data center, apply the PBR route maps to Core-Router1.
7. Determine which PBR method to use to verify PBR next-hop availability of a WAE. For more information, see the "Methods of Verifying PBR Next-Hop Availability" section on page 5-39.

Note For a description of the PBR commands that are referenced in this section, see the *Cisco Quality of Service Solutions Command Reference*.

Figure 5-5 shows that the WAEs (Edge-WAE1 and Core-WAE1) must reside in an out-of-band network that is separate from the traffic's destination and source. For example, Edge-WAE1 is on a subnet separate from the clients (the traffic source), and Core-WAE is on a subnet separate from the file servers and application servers (the traffic destination). Additionally, the WAE may need to be connected to the router that is redirecting traffic to it through a tertiary interface (a separate physical interface) or subinterface to avoid a routing loop. For more information on this topic, see the "Using Tertiary Interfaces or Subinterfaces to Connect WAEs to Routers" section on page 2-24.

Figure 5-5 Example of Using PBR or WCCP Version 2 for Transparent Redirection of All TCP Traffic to WAEs

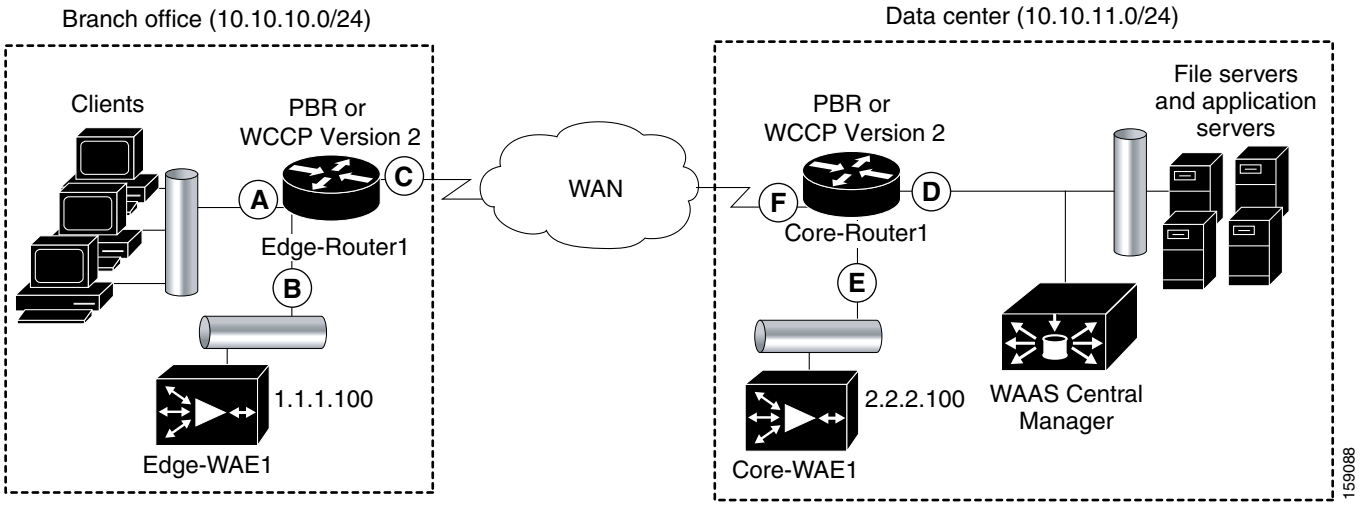

Table 5-3 provides a summary of the router interfaces that you must configure to use PBR or WCCP Version 2 to transparently redirect traffic to a WAE.

Table 5-3 Router Interfaces for WCCP or PBR Traffic Redirection to WAEs

Router interface	Comment
Edge-Router1	
A	Edge LAN interface (ingress interface) that performs redirection on outbound traffic.
B	Tertiary interface (separate physical interface) or a subinterface off of the LAN port on Edge-Router1. Used to attach Edge-WAE1 to Edge-Router1 in the branch office.
C	Edge WAN interface (egress interface) on Edge-Router1 that performs redirection on inbound traffic.
Core-Router1	
D	Core LAN interface (ingress interface) that performs redirection on outbound traffic.
E	Tertiary interface or subinterface off of the LAN port on Core-Router1. Used to attach Core-WAE1 to Core-Router1 in the data center.
F	Core WAN interface (egress interface) on Core-Router1 that performs redirection on inbound traffic.

Note In Figure 5-5, redundancy (for example, redundant routers, switches, WAEs, WAAS Central Managers, and routers) is not depicted.

The example in the "Configuring Policy-Based Routing" section on page 5-36 shows how to configure PBR as the traffic redirection method in a WAAS network that has one WAE in a branch office and one WAE in the data center (as shown in Figure 5-5).

Note The commands that are used to configure PBR on a router, can vary based on the Cisco IOS release installed on the router. For information about the commands that are used to configure PBR for the Cisco IOS release that you are running on your routers, see the appropriate Cisco IOS configuration guide.

Configuring Policy-Based Routing

The example in this section shows how to configure PBR as the traffic redirection method in a WAAS network that has one WAE in a branch office and one WAE in the data center (as shown in Figure 5-5).

To configure PBR to transparently redirect TCP traffic to WAEs, follow these steps:

Step 1 In the branch office, use extended IP access lists to specify which traffic is of interest to the LAN interface (ingress interface-A) on Edge-Router:

a. On Edge-Router1, define an extended IP access list within the range of 100 to 199. For example, create access list 100 on Edge-Router1:

```
Edge-Router1(config)# ip access-list extended 100
```

b. On Edge-Router1, specify which traffic is of interest to this particular interface:

- For example, mark any IP/TCP traffic from any local source addresses (traffic for any branch office clients) on any TCP port to any destination as interesting:

```
Edge-Router1(config-ext-nac1)# permit tcp 10.10.10.0 0.0.0.255 any
```

- Alternatively, you can selectively mark interesting traffic by defining the source IP subnet, destination IP address, and TCP port numbers. For example, mark IP/TCP traffic from any local source address on TCP ports 135 and 80 to any destination as interesting:

```
Edge-Router1(config-ext-nac1)# permit tcp 10.10.10.0 0.0.0.255 any eq 135
Edge-Router1(config-ext-nac1)# permit tcp 10.10.10.0 0.0.0.255 any eq 80
```

Step 2 In the branch office, use extended IP access lists to specify which traffic is of interest to the WAN interface (egress interface-C) on Edge-Router1:

a. On Edge-Router1, define an extended IP access list within the range of 100 to 199. For example, create access list 101 on Edge-Router1:

```
Edge-Router1(config)# ip access-list extended 101
```

b. On Edge-Router1, specify which traffic is of interest to its WAN interface:

- For example, mark any IP/TCP traffic to a local device as interesting:

```
Edge-Router1(config-ext-nac1)# permit tcp any 10.10.10.0 0.0.0.255
```

- Alternatively, you can selectively mark interesting traffic by defining the source IP subnet, destination IP address, and TCP port numbers. For example, mark IP/TCP traffic to any local source addresses on TCP ports 135 and 80 to any destination as interesting:

```
Edge-Router1(config-ext-nac1)# permit tcp any 10.10.10.0 0.0.0.255 eq 135
Edge-Router1(config-ext-nac1)# permit tcp any 10.10.10.0 0.0.0.255 eq 80
```

Step 3 In the data center, use extended IP access lists to specify which traffic is of interest to the LAN interface (ingress interface-D) on Core-Router1:

a. On Core-Router1, define an extended IP access list within the range of 100 to 199. For example, create access list 102 on Core-Router1:

```
Core-Router1(config)# ip access-list extended 102
```

b. On Core-Router1, specify which traffic is of interest to its LAN interface:

- For example, mark any IP/TCP traffic sourced from any local device (for example, traffic sourced from any file server or application server in the data center) on any TCP port to any destination as interesting:

```
Core-Router1(config-ext-nac1)# permit tcp 10.10.11.0 0.0.0.255 any
```

- Alternatively, you can selectively mark traffic as interesting by defining the source IP subnet, destination IP address, and TCP port numbers. For example, selectively mark IP/TCP traffic sourced from any local device on TCP ports 135 and 80 to any destination as interesting:

```
Core-Router1(config-ext-nac1)# permit tcp 10.10.11.0 0.0.0.255 any eq 135
Core-Router1(config-ext-nac1)# permit tcp 10.10.11.0 0.0.0.255 any eq 80
```

Step 4 In the data center, use extended IP access lists to mark traffic of interest for the WAN interface (egress interface-F) on Core-Router1:

a. On Core-Router1, define an extended access list within the range of 100 to 199. For example, create access list 103 on Core-Router1:

```
Core-Router1(config)# ip access-list extended 103
```

b. On Core-Router1, mark interesting traffic for the WAN interface:

- For example, mark any IP/TCP traffic destined to any local device (for example, traffic destined to any file server or application server in the data center) as interesting:

```
Core-Router1(config-ext-nac1)# permit tcp any 10.10.11.0 0.0.0.255
```

- Alternatively, you can selectively mark traffic as interesting by defining the source IP subnet, destination IP address, and TCP port numbers. For example, mark IP/TCP traffic on ports 135 and 80 to any local source addresses as interesting:

```
Core-Router1(config-ext-nac1)# permit tcp any 10.10.11.0 0.0.0.255 eq 135
Core-Router1(config-ext-nac1)# permit tcp any 10.10.11.0 0.0.0.255 eq 80
```

Step 5 In the branch office, define PBR route maps on Edge-Router1:

a. Define a route map for the LAN interface (ingress interface). In the following example, the WAAS-EDGE-LAN route map is created:

```
Edge-Router1(config)# route-map WAAS-EDGE-LAN permit
```

b. Define a route map for the WAN interface (egress interface).

In the following example, the WAAS-EDGE-WAN route map is created:

```
Edge-Router1(config)# route-map WAAS-EDGE-WAN permit
```

c. Specify the match criteria.

Use the **match** command to specify the extended IP access list that Edge-Router1 should use to determine which traffic is of interest to its WAN interface. If you do not specify a **match** command, the route map applies to all packets.

In the following example, Edge-Router1 is configured to use the access list 101 as the criteria for determining which traffic is of interest to its WAN interface:

```
Edge-Router1(config-route-map)# match ip address 101
```

Note The **ip address** command option matches the source or destination IP address that is permitted by one or more standard or extended access lists.

d. Specify how the matched traffic should be handled.

In the following example, Edge-Router1 is configured to send the packets that match the specified criteria to the next hop, which is Edge-WAE1 that has an IP address of 1.1.1.100:

```
Edge-Router1(config-route-map)# set ip next-hop 1.1.1.100
```

Note If you have more than one branch WAE, you can specify the IP address of a second branch WAE for failover purposes (for example, enter the **set ip next-hop 1.1.1.101** command on Edge-Router1) to specify a next-hop address of 1.1.1.101 (the IP address of Edge-WAE2) for failover purposes. The **next-hop** command is used for failover purposes and not for load-balancing purposes.

Step 6 In the data center, create route maps on Core-Router1:

a. Define a route map on the LAN interface (ingress interface).

In the following example, the WAAS-CORE-LAN route map is created:

```
Core-Router1(config)# route-map WAAS-CORE-LAN permit
```

b. Define a route map on the WAN interface (egress interface).

In the following example, the WAAS-CORE-WAN route map is created:

```
Core-Router1(config)# route-map WAAS-CORE-WAN permit
```

c. Specify the match criteria.

Use the **match** command to specify the extended IP access list that Core-Router 1 should use to determine which traffic is of interest to its WAN interface. If you do not enter a **match** command, the route map applies to all packets. In the following example, Core-Router1 is configured to use the access list 103 as the criteria for determining which traffic is of interest to its WAN interface:

```
Core-Router1(config-route-map)# match ip address 103
```

d. Specify how the matched traffic is to be handled.

In the following example, Core-Router1 is configured to send packets that match the specified criteria to the next hop, which is Core-WAE1 that has an IP address of 2.2.2.100:

```
Core-Router1(config-route-map)# set ip next-hop 2.2.2.100
```

Note If you have more than one data center WAE, you can specify the IP address of a second data center WAE for failover purposes (for example, enter the **set ip next-hop 2.2.2.101** command on Core-Router1) to specify a next-hop address of 2.2.2.101 (the IP address of Core-WAE2) for failover purposes. The **next-hop** command is used for failover purposes and not for load-balancing purposes.

Step 7 In the branch office, apply the route maps to the LAN interface (ingress interface) and the WAN interface (egress interface) on Edge-Router1:

a. On Edge-Router1, enter interface configuration mode:

```
Edge-Router1(config)# interface FastEthernet0/0.10
```

b. Specify that the LAN router interface should use the WAAS-EDGE-LAN route map for PBR:

```
Edge-Router1(config-if)# ip policy route-map WAAS-EDGE-LAN
```

c. Enter interface configuration mode:

```
Edge-Router1(config-if)# interface Serial0
```

d. Specify that the WAN router interface should use the WAAS-EDGE-WAN route map for PBR:

```
Edge-Router1(config-if)# ip policy route-map WAAS-EDGE-WAN
```

Step 8 In the data center, apply the route maps to the LAN interface (ingress interface) and the WAN interface (egress interface) on Core-Router1:

a. On Core-Router1, enter interface configuration mode:

```
Core-Router1(config)# interface FastEthernet0/0.10
```

b. Specify that for PBR, the LAN router interface should use the WAAS-CORE-LAN route map:

```
Core-Router1(config-if)# ip policy route-map WAAS-CORE-LAN
```

c. Enter interface configuration mode:

```
Core-Router1(config-if)# interface Serial0
```

d. Specify that for PBR, the WAN router interface should use the WAAS-CORE-WAN route map:

```
Core-Router1(config-if)# ip policy route-map WAAS-CORE-WAN
```

Methods of Verifying PBR Next-Hop Availability

When using PBR to transparently redirect traffic to WAEs, we recommend that you use one of the following methods to verify the PBR next-hop availability of a WAE. The method that you choose is based on the version of the Cisco IOS software that is running on the routers and the placement of your WAEs. However, method 2 is the preferred method whenever possible:

- Method 1—If the device sees the WAEs as a CDP neighbor (directly connected), it can use CDP and ICMP to verify that the WAE is operational. For more information, see the "Method 1: Using CDP to Verify Operability of WAEs" section on page 5-40.
- Method 2 (Recommended method)—If the device is running the Cisco IOS software Release 12.4 or later and the device does not see the WAE as a CDP neighbor, IP service level agreements (SLAs) can be used to verify that the WAE is operational using ICMP echoes. For more information, see the "Method 2: Using IP SLAs to Verify WAE Operability Using ICMP Echo Verification (Recommended Method)" section on page 5-40.
- Method 3—If the device is running the Cisco IOS software Release 12.4 or later and does not see the WAE as a CDP neighbor, IP SLAs can be used to verify that the WAE is operational using TCP connection attempts. For more information, see the "Method 3: Using IP SLAs to Verify WAE Operability Using TCP Connection Attempts" section on page 5-41.

Note In this section, device is used to refer to the router or switch that has been configured to use PBR to transparently redirect traffic to a WAE.

To verify whether the WAE is CDP visible to a device that has been configured to use PBR, enter the **show cdp neighbors** command on the device. If the WAE is CDP visible to the device, the WAE will be listed in the output of the **show cdp neighbors** command.

Method 1: Using CDP to Verify Operability of WAEs

If the device that is configured to use PBR views the WAEs as a CDP neighbor (the WAE is directly connected to the device), you can configure CDP and ICMP to verify the availability of a WAE as a PBR next hop.

The following example shows how to use this method to verify PBR next-hop availability of a WAE. You must complete the following configuration process for each of the LAN and WAN route maps that are configured when CDP should be used.

To use CDP to verify operability of WAEs, follow these steps:

Step 1 On the router where PBR is configured (for example, on the branch office router named Edge-Router1), enter configuration mode and enable CDP on the router:

```
Edge-Router1(config)# cdp run
```

Step 2 Enable route-map configuration mode for the route map, WAAS-EGDE-LAN, which has already been created on the router:

```
Edge-Router1(config)# route-map WAAS-EDGE-LAN permit
```

Step 3 Configure the router to use CDP to verify the availability of the configured next-hop addresses:

```
Edge-Router1(config-route-map)# set ip next-hop verify-availability
```

Step 4 Enable CDP on the WAE (for example, on the branch office WAE named Edge-WAE1) that you want the router to redirect traffic to using PBR:

```
Edge-WAE1(config)# cdp enable
```

If you are configuring PBR and have multiple WAEs and are using Method 1 to verify the PBR next-hop availability of a WAE, no additional configuration is necessary after you have completed the preceding process.

Method 2: Using IP SLAs to Verify WAE Operability Using ICMP Echo Verification (Recommended Method)

To use IP SLAs and ICMP (the recommended method) to verify PBR next-hop availability of a WAE, follow these steps:

Step 1 On the branch office router named Edge-Router1, enter the route-map configuration mode for the route map named WAAS-EDGE-LAN, which has been previously configured on this router:

```
Edge-Router1(config)# route-map WAAS-EDGE-LAN permit
```

Step 2 Specify a match condition for the traffic. In the following example, the match condition specifies access list number 105:

```
Edge-Router1(config)# match ip address 105
```

Step 3 Configure the route map to use IP SLA tracking instance number 1 to verify the availability of the next-hop WAE (for example, the branch WAE named Edge-WAE1 that has an IP address of 1.1.1.100):

```
Edge-Router1(config-route-map)# set ip next-hop verify-availability 1.1.1.100 track 1
```

Note Enter the **set ip next-hop verify-availability** command for each route-map that has been configured on this branch office edge router and on the data center's core router that has also been configured to use PBR to redirect traffic to WAEs.

Step 4 Configure the IP SLA tracking instance 1:

```
Edge-Router1(config-route-map)# exit
Edge-Router1(config)# ip sla 1
Edge-Router1(config-ip-sla)#
```

Step 5 Configure the router to echo Edge-WAE1 using the specified source interface:

```
Edge-Router1(config-ip-sla)# icmp-echo 1.1.1.100 source-interface FastEthernet 0/0.20
```

Step 6 Configure the router to perform the echo every 20 seconds:

```
Edge-Router1(config-ip-sla)# frequency 20
Edge-Router1(config-ip-sla)# exit
```

Step 7 Schedule the IP SLA tracking instance 1 to start immediately and to run continuously:

```
Edge-Router1(config)# ip sla schedule 1 life forever start-time now
```

Step 8 Configure the IP SLA tracking instance 1 to track the device, which is defined in the IP SLA tracking instance 1:

```
Edge-Router1(config)# track 1 rtr 1
```

If you are configuring PBR and have multiple WAEs, and you are using Method 2 to verify PBR next-hop availability of a WAE, you must configure a separate IP SLA per WAE and then run the **track** command per IP SLA.

Method 3: Using IP SLAs to Verify WAE Operability Using TCP Connection Attempts

If the device that is configured for PBR is running the Cisco IOS software Release 12.4 or later and does not see the WAE as a CDP neighbor, IP SLAs can be used to verify that the WAE is alive using TCP connection attempts. IP SLAs can be used to monitor a WAE's availability as the PBR next hop using TCP connection attempts at a fixed interval of 60 seconds.

To verify PBR next-hop availability of a WAE, follow these steps:

Step 1 On the branch office router named Edge-Router1, enter route-map configuration mode for the route map named WAAS-EDGE-LAN, which has been previously configured on this router:

```
Edge-Router1(config)# route-map WAAS-EDGE-LAN permit
```

Step 2 Configure the route map to use IP SLA tracking instance number 1 to verify the availability of the next-hop WAE (the Edge WAE that has an IP address of 1.1.1.100):

```
Edge-Router1(config-route-map)# set ip next-hop verify-availability 1.1.1.100 track 1
```

Note Enter the **set ip next-hop verify-availability** command for each route map that is configured on this branch office edge router and on the data center's core router that has also been configured to use PBR to transparently redirect traffic to WAEs.

Step 3 Configure the IP SLA tracking instance 1:

```
Edge-Router1(config-route-map)# exit
Edge-Router1(config)# ip sla 1
```

Step 4 Configure the router to use the specified source and destination ports to use TCP connection attempts at a fixed interval of 60 seconds to monitor the WAE availability:

```
Edge-Router1(config-ip-sla)# tcp-connect 1.1.1.100 80 source-port 51883 control disable
Edge-Router1(config-ip-sla)# exit
```

Step 5 Schedule the IP SLA tracking instance 1 to start immediately and to run forever:

```
Edge-Router1(config)# ip sla schedule 1 life forever start-time now
```

Step 6 Configure the IP SLA tracking instance 1 to track the device, which is defined in the IP SLA tracking instance 1:

```
Edge-Router1(config)# track 1 rtr 1
```

If you are configuring PBR and have multiple WAEs, and you are using Method 3 to verify PBR next-hop availability of a WAE, you must configure a separate IP SLA per WAE and then run the **track** command per IP SLA.

Using Inline Mode Interception

This section contains the following topics:

Information About Inline Interception

The WAE can physically and transparently intercept traffic between the clients and the router by using inline mode. To use inline mode, you must use a WAE with the Cisco WAE Inline Network Adapter or Interface Module installed. In this mode, you physically position the WAE device in the path of the traffic that you want to optimize, typically between a switch and a router, as shown in Figure 5-6. Redirection of traffic is not necessary.

Note When you install any inline WAE device, you must follow the cabling requirements described in the "Cabling" section of *Installing the Cisco WAE Inline Network Adapter* or the appropriate platform hardware guide.

Any combination of traffic interception mechanisms on peer WAEs is supported. For example, you can use inline interception on the branch WAE and WCCP on the data center WAE. For complex data center deployments, we recommend that you use hardware-accelerated WCCP interception with the WAAS AppNav solution (see Chapter 4, "Configuring AppNav") or load balancing with the Cisco Application Control Engine (ACE).

Figure 5-6 Inline Interception

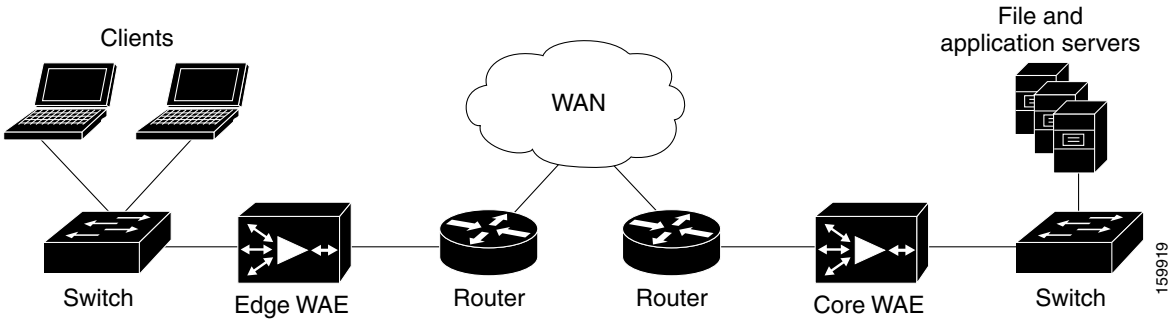

Note Inline mode and WCCP redirection are exclusive. You cannot configure inline mode if the WAE is configured for WCCP operation. Inline mode is the default mode when a Cisco WAE Inline Network Adapter is installed in a WAE device, but you must configure inline mode explicitly on a device with a Cisco Interface Module.

Note An inline WAE can be configured as a Central Manager, but the inline interception functionality is not be available.

The Cisco WAE Inline Network Adapter contains two or four Ethernet ports, the Cisco Interface Module contains two to eight Ethernet ports, and the Cisco AppNav Controller Interface Module contains four to 12 Ethernet ports. Ports on the Cisco WAE Inline Network Adapter are always configured as inline ports, while ports on the Interface Modules are configured as normal standalone ports by default, and you must explicitly configure these ports as inline ports. Each pair of inline ports is grouped into a logical inline group.

Each inline group has one LAN-facing port and one WAN-facing port. Typically, you use just one inline group, and connect the LAN-facing port to a switch and the WAN-facing port to a router. On adapters or interface modules with additional ports, the additional groups of interfaces are provided if you are using a network topology where you need to connect the WAE to multiple routers. Traffic that enters on one interface in a group exits the device on another interface in the same group.

Hardware platform support for inline ports is as follows:

- WAVE-274/474—Support one installed two-port Cisco WAE Inline Network Adapter.
- WAVE-574—Supports one installed two-port or four-port Cisco WAE Inline Network Adapter.

- WAE-674/7341/7371—Support up to two installed four-port Cisco WAE Inline Network Adapters, providing a total of eight inline ports.
- WAVE-294—Supports one installed Cisco Interface Module with 2, 4, or 8 ports.
- WAVE-594/694/7541/7571/8541—Support one installed Cisco Interface Module with 2, 4, or 8 ports or a Cisco AppNav Controller Interface Module with 4 or 12 ports.

Note The two-port 10-Gigabit Cisco Interface Module cannot be used in inline mode. The four-port 10-Gigabit Cisco AppNav Controller Interface Module is supported only on the WAVE-594.

You have the option of assigning an IP address to an inline interface, but it is not required. For more information, see the "Configuring an IP Address on an Inline Interface" section on page 5-51.

Traffic that flows through an inline group is transparently intercepted for optimization. Traffic that does not need to be optimized is bridged across the LAN/WAN interfaces. If a power, hardware, or unrecoverable software failure occurs, the network adapter automatically begins operating in bypass mode (fail-close), where all traffic is mechanically bridged between the LAN and WAN interfaces in each group. The Cisco WAE Inline Network Adapter and Cisco Interface Module also operate in bypass mode when the WAE is powered off or starting up. Additionally, you can manually put an inline group into bypass mode.

Note AppNav Controller Interface Modules do not support automatic bypass mode to continue traffic flow in the event of a failure. For high availability, two or more AppNav Controller Interface Modules should be deployed in an AppNav cluster. For more information on using inline mode with the AppNav solution, see Chapter 4, "Configuring AppNav."

Inline mode is configured by default to accept all TCP traffic. If the network segment in which the WAE is inserted is carrying 802.1Q tagged (VLAN) traffic, initially traffic on all VLANs is accepted. Inline interception can be enabled or disabled for each VLAN. However, optimization policies cannot be customized based on the VLAN.

You can serially cluster WAE devices operating in inline mode to provide higher availability if a device fails. For details, see the "Information About Clustering Inline WAEs" section on page 5-53.

Note When a WAE inline group enters bypass mode, the switch and router ports to which it is connected may have to reinitialize, which may cause an interruption of several seconds in the traffic flow through the WAE.

If the WAE is deployed in a configuration where the creation of a loop is not possible (that is, if it is deployed in a standard fashion between a switch and a router), configure PortFast on the switch port to which the WAE is connected. PortFast allows the port to skip the first few stages of the Spanning Tree Algorithm (STA) and move more quickly into a packet forwarding mode.

Enabling Inline Operation on WAEs

This section describes how to enable and configure inline settings on WAEs configured as application accelerators and that are not part of an AppNav Cluster (WAEs that are part of an AppNav Cluster use only the appnav-controller interception method). If you want to configure the inline settings on WAEs configured as AppNav Controllers, see the "Configuring Inline Operation on ANCs" section on page 5-49.

On WAVE-294/594/694/7541/7571/8541 devices that use Cisco Interface Modules, the Interface Module ports are configured by default for normal standalone operation. If you want to use the device in inline mode, you must configure the ports for inline operation. Enabling inline mode configures all ports for inline operation and converts each pair of ports to an inline group.

On other WAE devices that use the Cisco WAE Inline Network Adapter, the ports on the adapter always operate in inline mode. You can use this configuration window to enable or disable VLAN ID connection checking, which is the only setting that appears for such WAE devices.

To enable inline operation and configure general settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*. (You cannot enable inline operation from device groups.)

Step 2 Choose **Configure > Interception > Interception Configuration**.

Note If you are configuring a device using a WAAS version earlier than 5.0, choose **Configure > Interception > Inline > General Settings** to configure inline general settings. The configuration window looks different but has similar settings.

The Interception Configuration window appears.

Step 3 From the Interception Method drop-down list, choose **Inline** to enable inline mode. The Interception Method drop-down list is not shown for devices using WAAS versions earlier than 5.0.

The screen refreshes with the inline settings. (See Figure 5-7.)

Figure 5-7 Inline Interception Settings Window

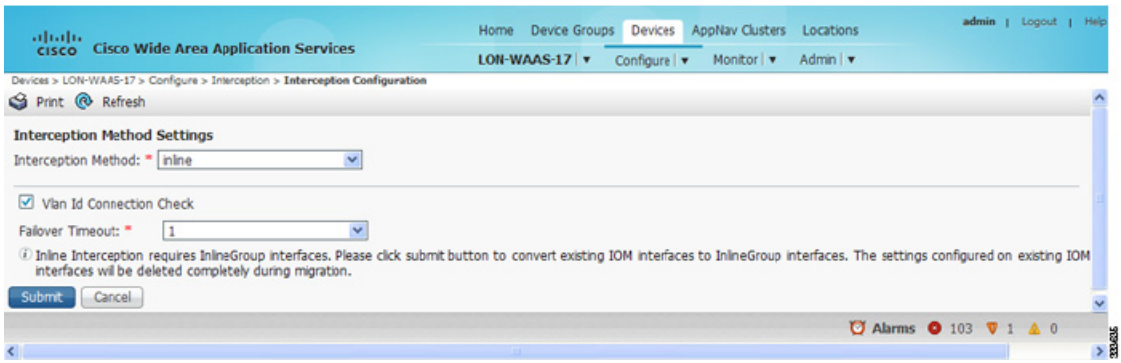

Step 4 Check the **Inline Enable** check box to enable inline operation.

The Inline Enable check box is shown only for WAVE devices using WAAS versions earlier than 5.0 and that have a Cisco Interface Module installed.

Step 5 Check the **Vlan ID Connection Check** check box to enable VLAN ID connection checking. Uncheck the check box to disable it. The default setting is enabled.

WAAS uses the VLAN ID to intercept or bridge VLAN traffic on the inline interface for a TCP flow. The VLAN ID of all packets sent in a particular TCP connection must match; any packets with a different VLAN ID will be bridged and not optimized. If your system has an asymmetric routing topology, in which the traffic flow in one direction uses a different VLAN ID than the traffic flow from the other direction, you may need to disable VLAN ID checking to ensure that the traffic is optimized.

Step 6 From the Failover Timeout drop-down list, choose the failover timeout (1, 5 or 25 seconds), which is the number of seconds that the interface should wait before going into bypass mode, after a device or power failure. The default is 1 second.

This item appears only for WAVE devices that use Cisco Interface Modules but not for AppNav Controller Interface Modules. For devices that use Cisco WAE Inline Network Adapters, the failover timeout is configured in the Inline Interface Settings window (see Figure 5-8 on page 5-47). This item is named Time Out for WAAS versions earlier than 5.0 and appears before the VLAN ID Connection Check item.

Step 7 Click **Submit**. A message appears for you to confirm that all Interface Module interfaces are to be converted to inline group interfaces and existing Interface Module interface configurations are to be removed.

Step 8 Click **OK** to confirm.

The inline groups are configured with basic default settings. To configure inline group settings, see the "Configuring Inline Interface Settings on WAEs" section on page 5-46.

For devices running WAAS versions earlier than 5.0, after enabling inline mode, it takes about two data feed poll cycles (about 10 minutes by default) for the inline groups to appear in the Inline Interfaces list in the lower part of the window.

Note Inline mode cannot be enabled if any of the Interface Module ports are configured as the primary interface. You must change the primary interface and return to this window to enable inline mode.

For devices running WAAS versions earlier than 5.0, if you configure any of the interfaces on a Interface Module with nondefault settings (standby group, port channel, BVI, speed, duplex, IP address, ACLs, and so on), inline mode cannot be enabled and a warning message appears that tells you to check all interfaces for any configuration settings. You must remove all configuration settings from all interface module interfaces (slot 1) and then return to this window to enable inline mode.

To enable inline operation from the CLI, use the **interception-method inline** global configuration command.

To configure VLAN ID checking from the CLI, use the **inline vlan-id-connection-check** global configuration command after inline operation is enabled.

Configuring Inline Interface Settings on WAEs

This section describes how to configure inline settings on WAEs configured as application accelerators and that are not part of an AppNav Cluster (WAEs that are part of an AppNav Cluster use only the appnav-controller interception method). If you want to configure the inline settings on WAEs configured as AppNav Controllers, see the "Configuring Inline Operation on ANCs" section on page 5-49.

To configure inline interface settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*. (You cannot configure inline interface settings from device groups.)

Step 2 Choose **Configure > Interception > Interception Configuration**.

Note If you are configuring a device using a WAAS version earlier than 5.0, choose **Configure > Interception > Inline > Inline Interfaces** to configure inline interface settings. The configuration window looks different but has similar settings.

The Inline Interfaces window appears, listing the inline interface groups available on the device.

Step 3 Choose an inline group to configure and click the **Edit** taskbar icon. (For devices using WAAS versions earlier than 5.0, click the **Edit** icon next to the interface.)

The Edit Inline Settings window appears, displaying the inline interface configurations for a particular slot and port group. (See Figure 5-8.)

Figure 5-8 Edit Inline Settings Window

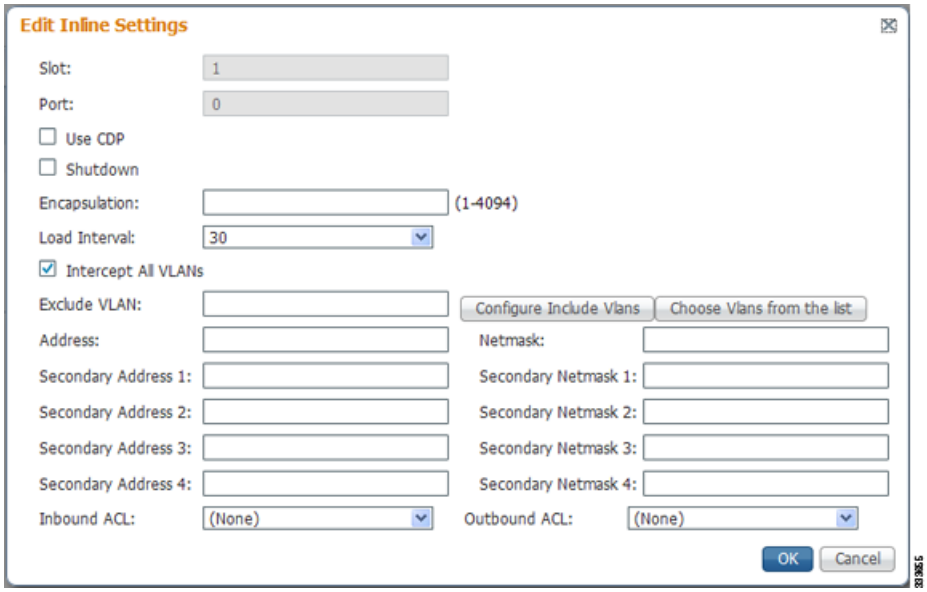

Step 4 Check the **Use CDP** check box to enable Cisco Discovery Protocol (CDP) on the inline group interfaces. The Use CDP check box is not shown for devices using WAAS versions earlier than 5.0.

When enabled, CDP obtains protocol addresses of neighboring devices and discovers the platform of those devices. It also shows information about the interfaces used by your router.

Configuring CDP from the CDP Settings window enables CDP globally on all the interfaces. For information on configuring CDP settings, see the "Configuring CDP Settings" section on page 6-25.

Step 5 Check the **Shutdown** check box to shut down the inline group. This setting bridges traffic across the LAN/WAN interfaces without any processing.

Step 6 In the Encapsulation field, enter the VLAN ID that is to be assigned to traffic that leaves the WAE. The VLAN ID should be set to match the VLAN ID expected by the router.

For more information about the VLAN ID, see the "Configuring an IP Address on an Inline Interface" section on page 5-51.

Step 7 From the Load Interval drop-down list, choose the interval in seconds at which to poll the interface for statistics and calculate throughput. The default is 30 seconds. (The Load Interval item is not shown for devices using WAAS versions earlier than 5.0.)

Step 8 Check the **Intercept all VLANs** check box to enable inline interception on the interface group. Inline interception is enabled by default when the WAE contains a Cisco WAE Inline Network Adapter but must be explicitly enabled on devices with a Cisco Interface Module (see the "Enabling Inline Operation on WAEs" section on page 5-44).

Step 9 In the Exclude VLAN field, enter a list of one or more VLAN ranges to exclude from optimization. You can enter the word "native" to exclude the native VLAN. Separate each VLAN range from the next with a comma. Alternatively, you can select VLAN ranges from a list by following these steps:

a. Click the **Configure Include VLANs** button when you know the list of VLANs that you want to include in inline interception. This button runs a script that prompts you for a comma-separated list of VLANs that you want to include. The script generates an inverse list of all VLANs that should be excluded and then updates the window and puts the list into the Exclude VLAN field.

b. Click the **Choose VLANs from the list** button to choose VLAN ranges. The VLAN Range Assignments window appears, displaying the VLAN ranges that are defined. Defining VLAN ranges is described in the "Configuring VLANs for Inline Support" section on page 5-52.

c. Choose the VLAN ranges to include or exclude by doing the following:

- Check the check box next to each VLAN range that you want to include for optimization and click the **Include Vlan** taskbar icon. All VLANs that are not included for optimization are excluded. For devices using WAAS versions earlier than 5.0, click next to each VLAN range that you want to include. The icon changes to .
- Check the check box next to each VLAN range that you want to exclude from optimization and click the **Exclude Vlan** taskbar icon. For devices using WAAS versions earlier than 5.0, click next to each VLAN range that you want to exclude from optimization. The icon changes to .
- Click the **Clear Selection** taskbar icon to clear all selections. For devices using WAAS versions earlier than 5.0, click in the taskbar to select all available VLAN ranges for optimization, or click in the taskbar to exclude all VLAN ranges from optimization.

d. Click **OK**. For devices using WAAS versions earlier than 5.0, click **Submit**.

Step 10 From the Failover Timeout drop-down list, choose **1**, **3**, **5**, or **10** seconds. The default is 1 second. This value sets the number of seconds after a failure event that the WAE waits before beginning to operate in bypass mode. In bypass mode, all traffic received on either port of the interface group is forwarded out the other port in the group.

This check box applies only to devices that use Cisco WAE Inline Network Adapters. For devices that use Cisco Interface Modules, the failover timeout is configured in the Inline Interception Settings window (see Figure 5-7 on page 5-45) and does not appear in this window.

Step 11 Configure the Speed and Mode port settings as follows (these settings are not used for interfaces on the Cisco Interface Module on a device using WAAS version 5.0 or later, which uses auto sensing):

a. Uncheck the **AutoSense** check box, which is enabled by default.

b. From the Speed drop-down list, choose a transmission speed (10, 100, 1000, or 10000 Mbps). You must choose 1000 Mbps for fiber Gigabit Ethernet interfaces on a Cisco Interface Module for devices using WAAS versions earlier than 5.0.

c. From the Mode drop-down list, choose a transmission mode (full-duplex or half-duplex). You must choose full-duplex for fiber Gigabit Ethernet interfaces on a Cisco Interface Module for devices using WAAS versions earlier than 5.0.

Note We strongly recommend that you do not use half-duplex connections on the WAE or on routers, switches, or other devices. Half duplex impedes performance and should not be used. Check each Cisco WAE interface and the port configuration on the adjacent device (router, switch, firewall, and WAE) to verify that full duplex is configured.

Step 12 In the Address field, enter an IP address for the inline interface, if you want to assign an IP address.

Step 13 In the Netmask field, enter a subnet mask for the inline interface.

Step 14 Enter up to four secondary IP addresses and corresponding subnet masks in the Secondary Address and Secondary Netmask fields.

Configuring multiple IP addresses allows the device to be present in more than one subnet and can be used to optimize response time because it allows the data to go directly from the WAAS device to the client that is requesting the information without being redirected through a router. The WAAS device becomes visible to the client because both are configured on the same subnet.

Step 15 In the Default Gateway field, enter the default gateway IP address. The Default Gateway field is not shown for devices using WAAS versions 5.0 or later.

Step 16 (Optional) From the Inbound ACL drop-down list, choose an IP ACL to apply to inbound packets.

The drop-down list contains all the IP ACLs that you configured in the system.

Step 17 (Optional) From the Outbound ACL drop-down list, choose an IP ACL to apply to outbound packets.

Step 18 Click **OK**. For devices using WAAS versions earlier than 5.0, click **Submit**.

Step 19 For WAAS version 5.0 and later, choose **Configure > Network > Default Gateway** to configure the default gateway for an inline interface.

a. In the Default Gateway field, enter the default gateway IP address.

b. Click **Submit**.

To configure inline interception from the CLI, use the **interface InlineGroup** global configuration command.

Configuring Inline Operation on ANCs

This section describes how to enable and configure inline settings on WAAS devices configured as AppNav Controllers (ANCs). You can also use the AppNav Cluster wizard to configure an inline ANC and create an inline bridge interface, as described in the "Creating a New AppNav Cluster with the Wizard" section on page 4-14.

If you want to configure the inline settings on WAEs configured as application accelerators, see the "Enabling Inline Operation on WAEs" section on page 5-44.

On WAVE-594/694/7541/7571/8541 devices that use Cisco AppNav Controller Interface Modules, the AppNav Controller Interface Module ports are configured by default for normal standalone operation. If you want to use the device in inline mode, you must configure the ports for inline operation and create an inline bridge group. Enabling inline mode configures all ports for inline operation.

To enable inline operation and configure an inline bridge group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*. (You cannot enable inline operation from device groups.)

Step 2 Choose **Configure > Interception > Interception Configuration**.

The Interception Configuration window appears.

Step 3 From the Interception Method drop-down list, choose **Inline** to enable inline mode.

Step 4 Click **Submit** to enable inline mode and refresh the window with additional settings.

All existing bridge groups are listed, showing the bridge group number, protocol, link state propagation setting, VLAN ranges, and included interfaces.

From this list, you can perform the following tasks:

- Edit the settings for a bridge group by choosing it and clicking the **Edit** taskbar icon.
- Delete a bridge group by choosing it and clicking the **Delete** taskbar icon.
- Create a new bridge group as described in the following steps.

Step 5 Click the **Create Bridge** taskbar icon.

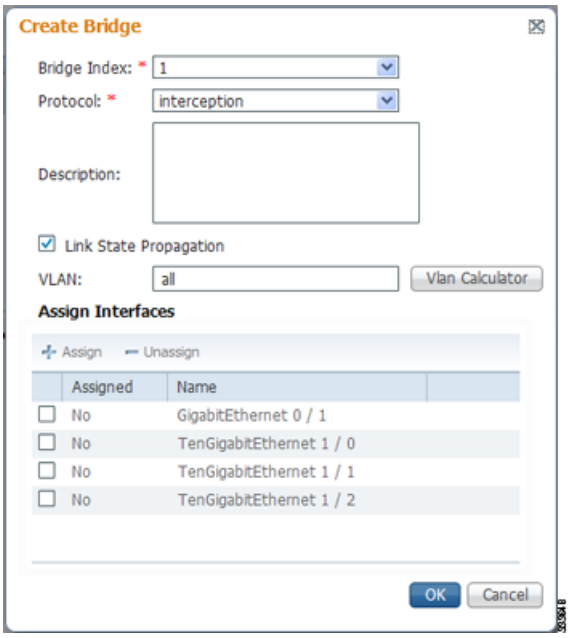

Step 6 From the Bridge Index drop-down list, choose the bridge group number.

Step 7 (Optional) In the Description field, enter a bridge group description.

Step 8 (Optional) Check the **Link State Propagation** check box to enable link state propagation. It is enabled by default.

Link state propagation means that if one interface in the inline bridge group is down, the system automatically shuts down the other interface to ensure that any network failover scheme is triggered.

Step 9 (Optional) Configure VLANs to include in interception. Initially all VLANS are included. If you want to include or exclude specific VLAN ranges, follow these steps:

a. Click the **Vlan Calculator** button.

b. For each VLAN range that you want to include in interception, set the **Select Operation Type** drop-down list to Add/Include. In the Vlan Range field, enter a comma-separated list of one or more VLAN ranges to include. You can enter the word "native" to include the native VLAN.

c. For each VLAN range that you want to exclude from interception, set the **Select Operation Type** drop-down list to Except/Exclude. In the Vlan Range field, enter a comma-separated list of one or more VLAN ranges to exclude. You can enter the word "native" to exclude the native VLAN.

d. Click **OK** to save your settings.

Step 10 In the Assign Interfaces area check the box next to two interfaces that you want to assign to this bridge group, then click the **Assign** taskbar icon. To unassign any assigned interfaces, check each interface that you want to unassign and click the **Unassign** taskbar icon. The bridge group can contain two physical or two port-channel interfaces, or a combination.

Step 11 Click **OK** to create the bridge group.

Configuring an IP Address on an Inline Interface

You can assign IP addresses to the inline group interfaces but it is not required. You can assign a primary IP address and up to four secondary IP addresses, using the procedure discussed in the "Configuring Inline Interface Settings on WAEs" section on page 5-46.

You can set an inline group interface as the primary interface on the WAE by using the Configure > Network > Network Interfaces window, in the Primary Interface drop-down list.

In scenarios where the primary interface for a WAE is set to an inline group interface and management traffic is configured on a separate IP address (either on a secondary IP address on the same inline group interface or on a built-in interface), you must configure the WAAS Central Manager to communicate with the WAE on the IP address designated for management traffic. Configure the WAE management interface settings with the Configure > Network > Management Interface Settings menu item. For WAAS versions earlier than 5.0, configure the WAE management traffic IP address in the *device-name* > Activation window, in the Management IP field.

If a WAE operating in inline mode is present in an 802.1Q VLAN trunk line between a switch and a router, and you are configuring the inline interface with an IP address, you must set the VLAN ID that is to be assigned to traffic that leaves the WAE. The VLAN ID should be set to match the VLAN ID expected by the router.

Use the **encapsulation dot1Q** interface command to assign a VLAN ID, as follows:

```
(config)# interface inlineGroup 1/0
(config-if)# encapsulation dot1Q 100
```

This example shows how to assign VLAN ID 100 to the traffic leaving the WAE. The VLAN ID can range from 1 through 4094.

Note You can set the VLAN ID of the inline traffic by using the **encapsulation dot1Q** interface command or by using the Central Manager menu item **Configure > Interception > Interception Configuration** (see the "Configuring Inline Interface Settings on WAEs" section on page 5-46).

If the VLAN ID that you set does not match the VLAN ID expected by the router subinterface, you may not be able to connect to the inline interface IP address.

The inline adapter supports only a single VLAN ID for each inline group interface. If you have configured a secondary address from a different subnet on an inline interface, you must have the same secondary address assigned on the router subinterface for the VLAN.

Using IEEE 802.1Q tunneling increases the frame size by 4 bytes when the tag is added. Therefore, you must configure all switches through which the tunneled packet traverses to be able to process larger frames by increasing the device MTU to at least 1504 bytes.

The following operating considerations apply to configuring IP addresses on the inline interfaces:

- This feature provides basic routable interface support and does not support the following additional features associated with the built-in interfaces: standby and port channel.
- If you have configured a WAE to use the inline interfaces for all traffic, inline interception must be enabled or the WAE will not receive any traffic.
- If you have configured a WAE to use the inline interfaces for all traffic and it goes into mechanical bypass mode, the WAE become inaccessible through the inline interface IP address. Console access is required for device management when an inline interface is in bypass mode.
- If you have configured a WAE with an IP address on an inline interface, the interface can accept only traffic addressed to it and ARP broadcasts, and the interface cannot accept multicast traffic.
- In a deployment using the Hot Standby Router Protocol (HSRP) where two routers that participate in an HSRP group are directly connected through two inline groups, HSRP works for all clients if the active router fails. However, this redundancy does not apply to the IP address of the WAE itself for management traffic, if management traffic is also configured to use the inline interface. If the active router fails, you will not be able to connect to the WAE inline IP address because the inline interface is physically connected to the failed router interface. You will be able to connect to the WAE through the second inline group interface that is connected to the standby router. If redundancy is needed for the IP address of the WAE itself for management traffic, we recommend that you use the IP addresses of the built-in interfaces rather than the inline interfaces.

Configuring VLANs for Inline Support

Initially, the WAE accepts traffic from all VLANs. You can configure the WAE to include or exclude traffic from certain VLANs; for excluded VLANs, traffic is bridged across the LAN/WAN interfaces in a group and is not processed.

To configure a VLAN for inline support, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Configure > Platform > Vlans**.

The Vlans window appears, which lists the VLANs that are defined. You can click the **Edit Vlan** icon next to an existing VLAN that you want to modify.

Step 2 In the taskbar, click the **Create New Vlan** icon. The Creating VLAN window appears.

Step 3 In the VLAN Name field, enter a name for the VLAN list.

Step 4 In the VLAN Ranges field, enter a list of one or more VLAN ranges. Separate each VLAN range from the next with a comma (but no space). This list of VLAN ranges can be included or excluded from optimization when you configure the inline interface group, as described in the "Configuring Inline Interface Settings on WAEs" section on page 5-46. You cannot specify the term "native" in this field.

Step 5 Click **Submit**.

This facility for creating VLAN lists is provided so that you can configure VLAN lists globally. You do not need to use this facility to configure VLANs for an inline interface. You can configure VLANs directly in the inline interface settings window, as described in the "Configuring Inline Interface Settings on WAEs" section on page 5-46.

Information About Clustering Inline WAEs

You can serially cluster two WAE devices that are operating in inline mode to provide higher availability in the data center if a device fails. If the current optimizing device fails, the inline group shuts down, or the device becomes overloaded, the second WAE device in the cluster provides the optimization services. Deploying WAE devices in a serial inline cluster for scaling or load balancing is not supported.

Note Overload failover occurs on TFO overload, not overload of individual application accelerators, and it is intended for temporary overload protection. We do not recommend that you continually run a WAE in an overloaded state, frequently triggering overload failover.

A serial cluster consists of two WAE devices connected together sequentially in the traffic path. The WAN port of one device is connected to the LAN port of the next device, as shown in Figure 5-9.

Figure 5-9 Inline Cluster

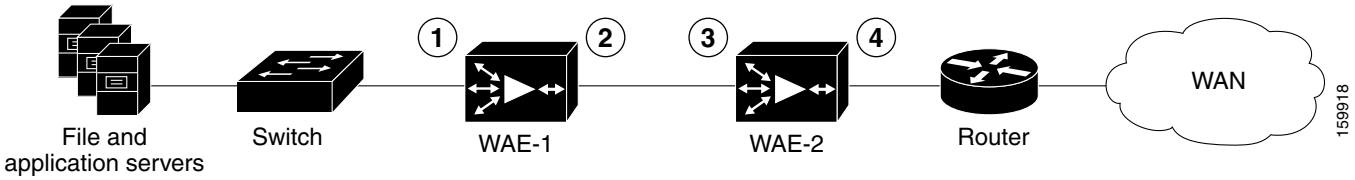

1	Inline LAN port on WAE-1	**3**	Inline LAN port on WAE-2
2	Inline WAN port on WAE-1	**4**	Inline WAN port on WAE-2

In a serial cluster, all traffic between the switch and router passes through all inline WAEs. In Figure 5-9, TCP connections are optimized by WAE-1. If WAE-1 fails, it bypasses the traffic and connections are then optimized by WAE-2.

The policy configuration of serially clustered WAEs should be the same. Additionally, we recommend that you use the same device for both WAEs in the cluster.

When serially clustering inline WAEs, on each WAE you must configure the address of the other WAE in the cluster as a non-optimizing peer. This disables optimization between the two peer WAEs in the serial cluster, since you want optimization only between the WAE peers on each side of the WAN link. For information on how to disable optimization between peers, see the "Disabling Peer Optimization Between Serial Inline WAEs" section on page 5-54.

Disabling Peer Optimization Between Serial Inline WAEs

To disable peer optimization between WAEs in a serial cluster, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*. (You cannot configure peer settings from device groups.)

Step 2 Choose **Configure > Peers > Peer Settings**.

The Peer Settings window appears. (See Figure 5-10.)

Figure 5-10 Peer Settings Window

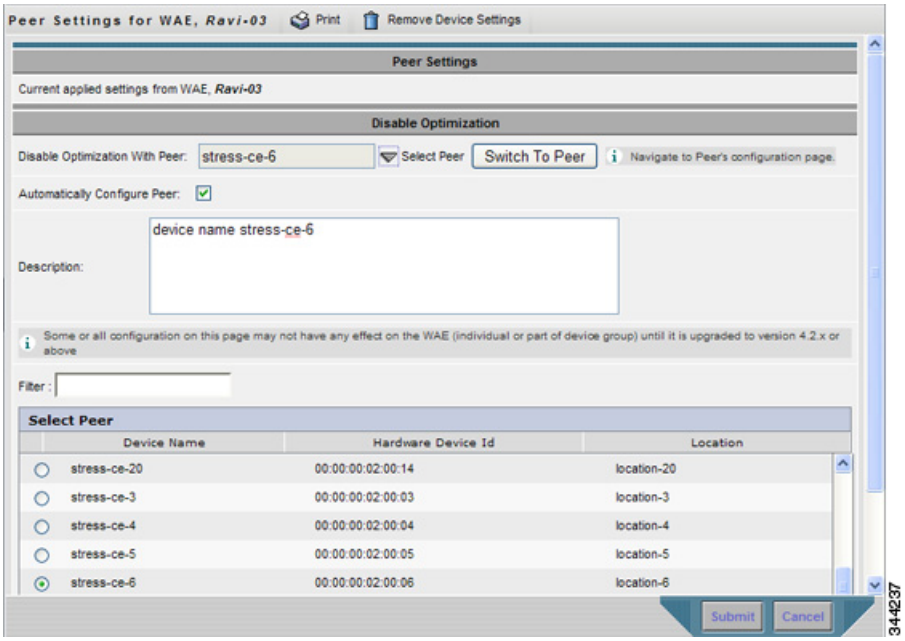

Step 3 Click the **Select Peer** triangle control to display in the lower part of the window other WAEs that are registered with this Central Manager (see the Select Peer area).

Step 4 In the Select Peer area, click the radio button next to the serial peer of the current device. The peer device name appears in the Disable Optimization With Peer field.

If you need to filter the device list, enter a string in the Filter field. As you enter characters, the device list is dynamically filtered to include only devices that have the filter string in their name or hardware device ID.

Step 5 Check the **Automatically Configure Peer** check box to allow the Central Manager to configure the other peer with a similar setting to disable optimization with the current device.

If you do not check this box, you must manually configure the other peer to disable optimization with the current device. After you submit your changes, you can click the **Switch to Peer** button to go to this same configuration page for the peer device.

Step 6 In the Description field, enter a description for the peer. The default description is the device name of the peer.

Step 7 Click **Submit**.

To disable serial peer optimization from the CLI, use the **no peer device-id** global configuration command. To reenable serial peer optimization, use the **peer device-id** global configuration command.

To view the status of all serial cluster pairs registered with the Central Manager, from the WAAS Central Manager menu, choose **Configure > Global > Peer Settings**. The Peer Settings status window appears, as shown in Figure 5-11.

Figure 5-11 Peer Settings For All Devices Window

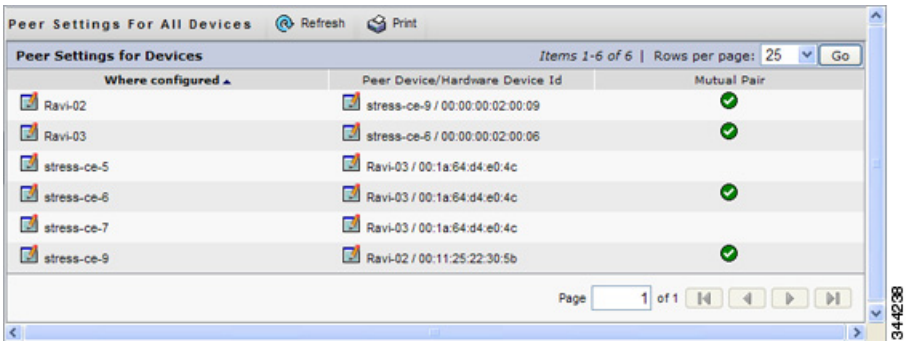

The window lists each WAE for which you have configured peer optimization settings. Verify that there are two entries for each serial cluster pair, both with a check mark in the Mutual Pair column. There should be an entry for each WAE in the pair (for example, the first and last entries in the figure).

If you see an entry without a check mark in the Mutual Pair column (like the third one in the figure), it indicates a WAE on which a serial peer is configured, but the peer is not similarly configured with the first device as its serial peer.

Configuring VPATH Interception on a vWAAS Device

VPATH intercepts traffic from the VM server, redirects it to a vWAAS device for WAN optimization, and then returns the response back to the Virtual Ethernet Module (VEM). The vWAAS egress traffic received by the VEM is forwarded without further VPATH interception.

Interception is configured on the server VM port profile in both directions.

To configure VPATH interception on a vWAAS device, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*. (You cannot configure vWAAS interface settings from device groups.)

Step 2 Choose **Configure > Interception > Interception Configuration**. The VPATH settings window appears.

Note If you are configuring a device using a WAAS version earlier than 5.0, choose **Configure > Interception > VPATH** to configure VPATH settings.

Step 3 From the Interception Method drop-down list, choose **vn-service** (**VPATH** on devices using WAAS versions earlier than 5.0) to enable VPATH interception on the vWAAS device.

Step 4 On devices using WAAS versions earlier than 5.0, check the **Enable VPATH** check box to enable VPATH interception on the vWAAS device. This check box is not editable on devices using WAAS versions 5.0 or later.

Note Only one type of interception can be enabled at a time.

Step 5 Click **Submit**.

To enable VPATH from the CLI, use the **interception-method vn-service vpath** global configuration command. The default is disabled. For monitoring and troubleshooting, use the **show statistics vn-service vpath** and **clear statistics vn-service vpath** EXEC configuration commands.

For more information on virtual WAAS configuration, see the *Cisco Wide Area Application Services vWAAS Installation and Configuration Guide*.

Configuring AppNav Interception

For WAEs that are part of an AppNav deployment and are configured as WAAS nodes (WNs) in an AppNav Cluster, you must configure them to use the appnav-controller interception method. These WNs receive traffic only from the ANCs, not directly from routers. It is on the ANC devices that you configure an interception method such as WCCP, PBR, or inline to intercept network traffic. For more information about an AppNav deployment, see Chapter 4, "Configuring AppNav."

If you create an AppNav Cluster by using the Central Manager wizard, or you add WNs to a cluster through the AppNav Clusters window, the Central Manager automatically configures WNs with the appnav-controller interception method. Once the WN is added to a cluster, its interception method cannot be changed.

To manually configure appnav-controller interception on a WN device, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Interception > Interception Configuration**. The Interception Configuration window appears.

Step 3 From the Interception Method drop-down list, choose **appnav-controller** to enable the appnav-controller interception method.

Step 4 Click **Submit**.

CHAPTER 6

Configuring Network Settings

This chapter describes how to configure basic network settings such as configuring additional network interfaces to support network traffic, creating port channel and standby interfaces, creating bridge interfaces for virtual blades, configuring optimization on WAAS Express interfaces, specifying a default gateway and DNS servers, enabling the Cisco Discovery Protocol (CDP), and configuring the directed mode of operation where peer WAEs exchange traffic using UDP encapsulation to avoid firewall traversal issues.

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

This chapter contains the following sections:

For information on configuring a bridge group for inline interfaces on an AppNav Controller Interface Module, see the "Configuring Inline Operation on ANCs" section on page 5-49 or use the AppNav Cluster wizard as described in the "Creating a New AppNav Cluster with the Wizard" section on page 4-14.

Configuring Network Interfaces

During initial setup, you chose an initial interface and either configured it for DHCP or gave it a static IP address, as described in the *Cisco Wide Area Application Services Quick Configuration Guide*. This section describes how to configure additional interfaces using options for redundancy, load balancing, and performance optimization.

This section contains the following topics:

We recommend that you use the WAAS Central Manager instead of the WAAS CLI to configure network settings, but if you want to use the CLI, see the following commands in the *Cisco Wide Area Application Services Command Reference*: **interface**, **ip address**, **port-channel**, and **primary-interface**.

Network interfaces are named as follows on WAAS devices:

- WAVE-274/474—Have one built-in Ethernet interface named GigabitEthernet 1/0.
- WAE-512/612/7326/674/7341/7371 and WAVE-574—Have two built-in Ethernet interfaces named GigabitEthernet 1/0 and GigabitEthernet 2/0.
- WAVE-294/594/694/7541/7571/8541—Have two built-in Ethernet interfaces named GigabitEthernet 0/0 and GigabitEthernet 0/1. Additional interfaces on the Cisco Interface Module and AppNav Controller Interface Module are named GigabitEthernet 1/0 to 1/11 or TenGigabitEthernet 1/0 to 1/3, depending on the number and type of ports.
- NME-WAE devices—Have an internal interface to the router that is designated 1/0 and an external interface that is designated 2/0.
- SM-SRE devices—Have an internal interface to the router that is designated 1/0 and an external interface that is designated 2/0.

Note We strongly recommend that you do not use half-duplex connections on the WAE or on routers, switches, or other devices. Half duplex impedes performance and should not be used. Check each Cisco WAE interface and the port configuration on the adjacent device (router, switch, firewall, and WAE) to verify that full duplex is configured.

When connecting an AppNav Controller to a Cisco Nexus 7000 Series switch, the interfaces on both devices must be set to the same auto-negotiate setting: either both on or both off. If they are set differently, switch link flapping can occur.

Note Layer 3 interfaces may drop bridge protocol data unit (BPDU) packets. This does not affect data traffic.

Configuring a Standby Interface

In this procedure, you configure a logical interface called a standby interface. After you configure this standby interface, you must associate physical or port-channel interfaces with the standby interface to create the standby group. In the WAAS Central Manager, you create the standby group by assigning two interfaces to the standby group and assigning one as primary.

Standby interfaces remain unused unless a member interface that is in use fails. When an in-use network interface fails (because of cable trouble, Layer 2 switch failure, or other failure), the other member interface of the standby group changes its state to in use and starts to carry traffic and take the load off the failed interface. With the standby interface configuration, only one interface is in use at a given time.

To configure standby interfaces, you must assign two physical or two port-channel interface members to a standby group. The following operating considerations apply to standby groups:

- A standby group consists of two physical or two port-channel interfaces. (If you are configuring a WAAS device running a version earlier than 5.0, both interfaces must be physical interfaces.)
- The maximum number of standby groups on a WAAS device is two. When using a Cisco AppNav Controller Interface Module, you can have up to three standby groups.
- A standby group is assigned a unique standby IP address, shared by all members of the group.
- Configuring the duplex and speed settings of the standby group member interfaces provides better reliability.
- IP ACLs can be configured on physical interfaces that are members of a standby group.
- One interface in a standby group is designated as the primary standby interface. Only the primary interface uses the group IP address.
- If the in-use interface fails, another interface in its standby group takes over and carries the traffic.
- If all the members of a standby group fail, then one recovers, the WAAS software brings up the standby group on the operational interface.
- The primary interface in a standby group can be changed at runtime. (The default action is to preempt the currently in-use interface if a different interface is made primary.)
- If a physical interface is a member of a standby group, it cannot also be a member of a port channel.
- If a device has only two interfaces, you cannot assign an IP address to both a standby group and a port channel. On such a device, only one logical interface can be configured with an IP address.
- The member interfaces of a standby group can be connected to different switches if you use a VLAN tagging protocol and assign the same VLAN tag to each interface.
- You cannot include a built-in Ethernet port and a port on a Cisco Interface Module in the same standby group.

Configuring a standby interface differs, depending on the version of the WAAS device that you are configuring. See one of the following topics:

- Configuring a Standby Interface on a Device with Version 5.0 or Later, page 6-4
- Configuring a Standby Interface on a Device Earlier than Version 5.0, page 6-5

Configuring a Standby Interface on a Device with Version 5.0 or Later

To configure a standby interface for devices with WAAS version 5.0 or later, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Network Interfaces**. The Network Interfaces window for the device appears. (See Figure 6-1.)

Figure 6-1 Network Interfaces for Device Window

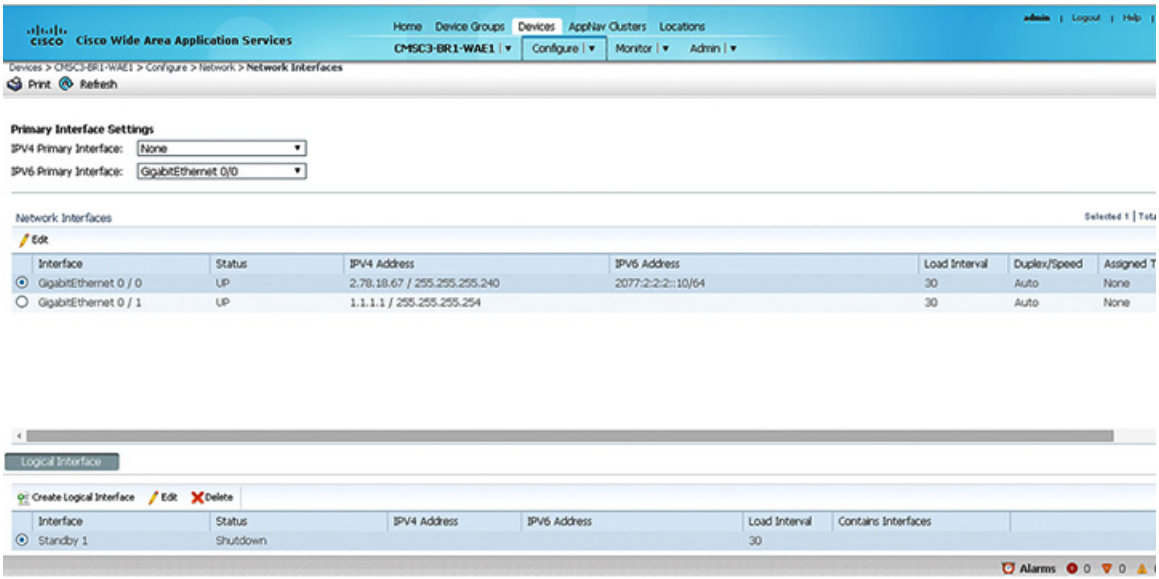

Step 3 In the taskbar of the lower area, click the **Create Logical Interface** icon. The Create Logical Interface window appears.

Step 4 From the Logical Interface Type drop-down list, choose **Standby** and click **OK**. The window refreshes with fields for configuring the standby group settings.

Step 5 From the Standby Group Number drop-down list, choose a group number for the interface.

Step 6 (Optional) From the Bridge Group Number drop-down list, choose a bridge virtual interface (BVI) group number with which to associate this standby interface, or **None**. For more information on BVI, see the "Bridging to a Virtual Blade Interface" section on page 6-17. This configuration item is not supported on AppNav Controller Interface Module ports.

Step 7 (Optional) In the Description field, enter a description for the standby group.

Step 8 (Optional) Check the **Shutdown** check box to shut down the hardware interface. By default, this option is disabled.

Step 9 (Optional) From the Load Interval drop-down list, choose the interval in seconds at which to poll the interface for statistics and calculate throughput. The default is 30 seconds.

Step 10 In the Address field, specify the IP address of the standby group.

Step 11 In the Netmask field, specify the netmask of the standby group.

Step 12 In the Assign Interfaces area, check the boxes next to the two interfaces that you want to assign to this standby group and click the **Assign** taskbar icon. To unassign any assigned interfaces, check each interface that you want to unassign and click the **Unassign** taskbar icon.

If you want to have two port-channel interfaces as members of the standby group, do not assign any interfaces here. When you create the port-channel interfaces, you assign the standby group number in that window.

Step 13 To assign one physical interface as the primary (active) interface in the standby group, ensure that it is the only interface checked and then click the **Enable Primary** taskbar icon.

Step 14 Click **OK**.

Configuring a Standby Interface on a Device Earlier than Version 5.0

To configure a standby interface for devices with WAAS versions earlier than 5.0, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Network Interfaces**. The Network Interfaces window for the device appears.

Step 3 In the taskbar, click the **Create New Interface** icon. The Creating New Network Interface window appears.

Step 4 From the Port Type drop-down list, choose **Standby**. The window refreshes with fields for configuring the standby group settings.

Step 5 From the Standby Group Number drop-down list, choose a group number for the interface.

Step 6 (Optional) In the Description field, optionally enter a description for the standby group.

Step 7 In the Address field, specify the IP address of the standby group.

Step 8 In the Netmask field, specify the netmask of the standby group.

Step 9 (Optional) Check the **Shutdown** check box to shut down the hardware interface. By default, this option is disabled.

Step 10 In the Default Gateway field, enter the default gateway IP address. If an interface is configured for DHCP, then this field is read only.

Step 11 (Optional) From the Bridge Group Number drop-down list, choose a bridge virtual interface (BVI) group number with which to associate this standby interface, or choose **None**. For more information on BVI, see the "Bridging to a Virtual Blade Interface" section on page 6-17.

Step 12 Click **Submit**.

Step 13 Configure the physical interface members as described in the "Assigning Physical Interfaces to the Standby Group" section on page 6-5.

After you create the standby interface, you need to assign two physical interfaces to the standby group.

Assigning Physical Interfaces to the Standby Group

After you have configured a logical standby interface for a device with a WAAS version earlier than 5.0, you configure the standby group by assigning physical interfaces to the standby group and setting one physical interface as the primary standby interface. The primary interface in the standby group uses the

standby group IP address. You must have a standby interface configured before you can set it as primary. (See the "Configuring a Standby Interface" section on page 6-3.)

You can assign an interface to a standby group only if the interface does not have an IP address assigned. The interface uses the IP address of the standby group.

Note Removing a physical interface from standby group 2 on all WAAS device models can cause network disruption for up to 30 seconds. Additionally, removing a physical interface from standby group 1 on device models WAE-612/674/7341/7371 and WAVE-574 can cause network disruption for up to 30 seconds. The best practice is to make such changes when traffic interception is disabled or at an off-peak time when traffic disruption is acceptable.

To associate an interface with a standby group and set it as the primary standby interface, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Network Interfaces**. The Network Interfaces window for the device appears.

Step 3 Click the **Edit** icon next to the physical interface that you want to assign to a standby group. The Interface Settings window appears.

Choose a physical interface, not a logical interface (standby, port channel, or BVI), in this step.

Step 4 Complete the following steps to assign the interface to a standby group and specify it as the primary standby interface:

a. In the Port Type To Assign drop-down list, choose **Standby**.

b. Check either the **Join Standby Group 1** or **Join Standby Group 2** check box. (Only one check box is shown if only one standby interface has been defined.)

c. (Optional) Check the **Standby Primary** check box if you want this physical interface to be the primary (active) interface in the standby group.

Step 5 Click **Submit**.

Configuring Multiple IP Addresses on a Single Interface

You can configure up to four secondary IP addresses on a single interface. This configuration allows the device to be present in more than one subnet and can be used to optimize the response time because it allows the data to go directly from the WAAS device to the client that is requesting the information without being redirected through a router. The WAAS device becomes visible to the client because both are configured on the same subnet.

Configuring multiple IP addresses is not supported on AppNav Controller Interface Module ports.

To configure multiple IP addresses on a single interface, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Network Interfaces**. The Network Interfaces listing window appears.

Step 3 Choose the physical interface that you want to modify and click the **Edit** taskbar icon. (For devices using WAAS versions earlier than 5.0, click the **Edit** icon next to the interface.)

The Interface Settings window appears.

Note Do not choose a standby or port-channel interface in this step. You cannot configure multiple IP addresses on these types of interfaces.

Step 4 In the Secondary Address and Secondary Netmask fields 1 through 4, enter up to four different IP addresses and secondary netmasks for the interface.

Step 5 Click **Submit**.

Modifying Ethernet Interface Settings

To modify the settings of a physical Ethernet interface, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Network Interfaces**.

The Network Interfaces window appears, listing the configured network interfaces.

Note On NME-WAE and SM-SRE devices, the internal interface to the router is designated slot 1, port 0 and the external interface is designated slot 2, port 0. For NME-WAE configuration details, see the document *Configuring Cisco WAAS Network Modules for Cisco Access Routers*. For SM-SRE configuration details, see the document *Cisco SRE Service Module Configuration and Installation Guide*.

Step 3 Choose the physical interface that you want to modify and click the **Edit** taskbar icon. (For devices using WAAS versions earlier than 5.0, click the **Edit** icon next to the interface.)

The Interface Settings window appears, displaying the interface configurations on a particular slot and port. The interface type, slot, and port are determined by the hardware.

Note When configuring the internal interface (GigabitEthernet 1/0) on an NME-WAE or SM-SRE device, you cannot change the following fields or check boxes: Port Channel Number, AutoSense, Speed, Mode, Address, Netmask, Use DHCP, and Standby Group. If you attempt to change these values, the Central Manager displays an error when you click Submit. These settings for the internal interface can be configured only through the host router CLI. For NME-WAE details, see the document *Configuring Cisco WAAS Network Modules for Cisco Access Routers*. For SM-SRE details, see the document *Cisco SRE Service Module Configuration and Installation Guide*.

Step 4 (Optional) In the Description field, enter a description for the interface.

Step 5 (Optional) Check the **Use CDP** check box to enable the Cisco Discovery Protocol (CDP) on an interface.

When enabled, CDP obtains protocol addresses of neighboring devices and discovers the platform of those devices. It also shows information about the interfaces used by your router.

Configuring CDP from the CDP Settings window enables CDP globally on all the interfaces. For information on configuring CDP settings, see the "Configuring CDP Settings" section on page 6-25.

Step 6 (Optional) Check the **Shutdown** check box to shut down the hardware interface.

Step 7 (Optional) From the Load Interval drop-down list, choose the interval in seconds at which to poll the interface for statistics and calculate throughput. The default is 30 seconds. (The Load Interval item is not shown for devices using WAAS versions earlier than 5.0.)

Step 8 (Optional) Check the **AutoSense** check box to set the interface to autonegotiate the speed and mode. (This setting is not available on interfaces on some Cisco Interface Modules.)

Checking this check box disables the manual Speed and Mode drop-down list settings.

Note When autosense is on, manual configurations are overridden. You must reboot the WAAS device to start autosensing.

Step 9 (Optional) Manually configure the interface transmission speed and mode settings as follows (these settings are not available on interfaces on some Cisco Interface Modules):

a. Uncheck the **AutoSense** check box.

b. From the Speed drop-down list, choose a transmission speed (10, 100, 1000, or 10000 Mbps). You must choose 1000 Mbps for fiber Gigabit Ethernet interfaces on a Cisco Interface Module.

c. From the Mode drop-down list, choose a transmission mode (**full-duplex** or **half-duplex**). You must choose full-duplex for fiber Gigabit Ethernet interfaces on a Cisco Interface Module. This configuration item is not supported on AppNav Controller Interface Module ports.

   Full-duplex transmission allows data to travel in both directions at the same time through an interface or a cable. A half-duplex setting ensures that data only travels in one direction at any given time. Although full duplex is faster, the interfaces sometimes cannot operate effectively in this mode. If you encounter excessive collisions or network errors, you may configure the interface for half-duplex rather than full duplex.

   Note We strongly recommend that you do not use half-duplex connections on the WAE or on routers, switches, or other devices. Half duplex impedes performance and should not be used. Check each Cisco WAE interface and the port configuration on the adjacent device (router, switch, firewall, and WAE) to verify that full duplex is configured.

Step 10 Specify a value (in bytes) in the MTU field to set the interface Maximum Transmission Unit (MTU) size.

The range is 576–1500 bytes. The MTU is the largest size of IP datagram that can be transferred using a specific data link connection.

Note The MTU field is not editable if the interface is assigned to a standby or port-channel group, or if a system jumbo MTU is configured.

Step 11 (Optional) Check the **Use DHCP** check box to obtain an interface IP address through DHCP. Checking this box hides the IP address and Netmask fields. (For devices with WAAS versions earlier than 5.0, these fields are not hidden but become grayed out.) This configuration item is not supported on AppNav Controller Interface Module ports.

Optionally supply a hostname in the Hostname field and a client ID in the Client Id field.

Step 12 In the Address field, enter a new IP address to change the interface IP address.

Step 13 In the Netmask field, enter a new netmask to change the interface netmask.

Step 14 (Optional) Enter up to four secondary IP addresses and corresponding subnet masks in the Secondary Address and Secondary Netmask fields. These fields are not supported on AppNav Controller Interface Module ports.

Configuring multiple IP addresses allows the device to be present in more than one subnet and can be used to optimize the response time because it allows the data to go directly from the WAAS device to the client that is requesting the information without being redirected through a router. The WAAS device becomes visible to the client because both are configured on the same subnet.

Step 15 In the Default Gateway field, enter the default gateway IP address. If an interface is configured for DHCP, then this field is read only. (The Default Gateway field is not shown for devices using WAAS versions 5.0 or later; instead configure it as described in the "Configuring the Default Gateway" section on page 6-9.)

Step 16 (Optional) From the Inbound ACL drop-down list, choose an IP ACL to apply to inbound packets.

The drop-down list contains all the IP ACLs that you configured in the system.

Step 17 (Optional) From the Outbound ACL drop-down list, choose an IP ACL to apply to outbound packets.

Step 18 Click **OK**. (For devices using WAAS versions earlier than 5.0, click **Submit**.)

Note Changing the interface transmission speed, duplex mode, or MTU can cause network disruption for up to 30 seconds. The best practice is to make such changes when traffic interception is disabled or at an off-peak time when traffic disruption is acceptable.

Configuring the Default Gateway

On WAAS devices with version 5.0 or later, configure the default gateway as follows:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Default Gateway**.

The Default Gateway window appears.

Step 3 In the Default Gateway field, enter the default gateway IP address.

Step 4 Click **Submit**.

To configure a default gateway from the CLI, you can use the **ip default-gateway** global configuration command.

On WAAS devices with versions earlier than 5.0, the default gateway is configured within the interface settings for each interface.

Configuring Port-Channel Settings

The WAAS software supports the grouping of up to four (eight on AppNav Controller Interface Modules) physical network interfaces into one logical interface called a port channel. After you configure this port-channel interface, you must associate physical interfaces with the port channel.

You can configure up to four port-channel interfaces (seven on AppNav Controller Interface Modules). This capability also provides interoperability with Cisco routers, switches, and other networking devices or hosts supporting EtherChannel, load balancing, and automatic failure detection and recovery based on each interface's current link status. EtherChannel is also referred to as a port channel.

You can use a port channel in standby interface, a bridge virtual interface (BVI) for a virtual blade, or as a member of an inline bridge group on an AppNav Controller Interface Module. For more information on configuring a BVI, see the "Bridging to a Virtual Blade Interface" section on page 6-17. For more information on configuring a bridge group on an AppNav Controller Interface Module, see the "Configuring Inline Operation on ANCs" section on page 5-49 or use the AppNav Cluster wizard as described in the "Creating a New AppNav Cluster with the Wizard" section on page 4-14.

The following operating considerations apply to a port-channel virtual interface:

- A physical interface can be a member of a port channel or a standby group, but not both.
- You cannot assign an IP address to both a port channel and a standby group. Only one logical interface can be configured with an IP address.
- All port-channel member interfaces must have the same port bandwidth.
- Port-channel settings are not applicable to vWAAS devices.
- You cannot include a built-in Ethernet port and a port on a Cisco Interface Module in the same port-channel interface.

Note You must disable autoregistration if the device has only two interfaces and both device interfaces are configured as port-channel interfaces.

Configuring a port-channel interface differs, depending on the version of the WAAS device that you are configuring. See one of the following topics:

- Configuring a Port-Channel Interface on a Device with Version 5.0 or Later, page 6-10
- Configuring a Port-Channel Interface on a Device Earlier than Version 5.0, page 6-11

Configuring a Port-Channel Interface on a Device with Version 5.0 or Later

To configure a port-channel interface for devices with WAAS version 5.0 or later, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Network Interfaces**. The Network Interfaces window for the device appears.

Step 3 In the taskbar of the lower area, click the **Create Logical Interface** icon. The Create Logical Interface window appears.

Step 4 From the Logical Interface Type drop-down list, choose **PortChannel** and click **OK**. The window refreshes with fields for configuring the port-channel interface settings.

Step 5 From the Port Channel Number drop-down list, choose a number for the interface.

Step 6 (Optional) From the Bridge Group Number drop-down list, choose a bridge group number with which to associate this interface, or choose **None**. The bridge group number can be associated with a BVI or an inline bridge group defined on an AppNav Controller.

Step 7 (Optional) From the Standby Group Number drop-down list, choose a standby group number with which to associate this interface, or choose **None**.

You must create the standby group with no assigned interfaces before it appears as a choice in this list.

Step 8 (Optional) In the Description field, optionally enter a description for the interface.

Step 9 (Optional) Check the **Shutdown** check box to shut down the hardware interface. By default, this option is disabled.

If you plan to assign this port-channel interface to a standby interface, check this box.

Step 10 (Optional) From the Load Interval drop-down list, choose the interval in seconds at which to poll the interface for statistics and calculate throughput. The default is 30 seconds.

Step 11 In the Address field, specify the IP address of the interface.

If you are assigning this port-channel interface to a standby group, do not configure an IP address or netmask. The standby group supplies the IP address and netmask.

Step 12 In the Netmask field, specify the netmask of the interface.

Step 13 (Optional) From the Inbound ACL drop-down list, choose an IP ACL to apply to inbound packets.

The drop-down list contains all the IP ACLs that you configured in the system.

Step 14 (Optional) From the Outbound ACL drop-down list, choose an IP ACL to apply to outbound packets.

Step 15 In the Assign Interfaces area, click the check box next to the interfaces that you want to assign to this port channel and click the **Assign** taskbar icon. To unassign any assigned interfaces, check each interface that you want to unassign and click the **Unassign** taskbar icon.

If you plan to assign this port-channel interface to a standby interface, do not assign interfaces until after the port channel is assigned to the standby interface.

Step 16 Click **OK**.

Configuring a Port-Channel Interface on a Device Earlier than Version 5.0

To configure a port-channel interface for devices with WAAS versions earlier than 5.0, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Network Interfaces**. The Network Interfaces window appears, listing all the interfaces for the chosen device.

Step 3 In the taskbar, click the **Create New Interface** icon. The Creating New Network Interface window appears.

Step 4 From the Port Type drop-down list, choose **PortChannel**.

The window refreshes and provides fields for configuring the network interface settings.

Step 5 In the Port Channel Number drop-down list, choose the number of the port-channel interface. Up to four port channels are supported, depending on the WAAS device model and installed interface module.

Step 6 (Optional) In the Bridge Group Number drop-down list, choose the number of the bridge group to which you want to assign this port-channel interface, if you want to bridge to a virtual blade.

Step 7 (Optional) In the Description field, optionally enter a description for the port channel.

Step 8 (Optional) Check the **Shutdown** check box to shut down this interface. By default, this option is disabled.

Step 9 In the Default Gateway field, enter the default gateway IP address.

Step 10 In the Address field, specify the IP address of the interface.

Step 11 In the Netmask field, specify the netmask of the interface.

Step 12 (Optional) From the Inbound ACL drop-down list, choose an IP ACL to apply to inbound packets.

The drop-down list contains all the IP ACLs that you configured in the system.

Step 13 (Optional) From the Outbound ACL drop-down list, choose an IP ACL to apply to outbound packets.

Step 14 Click **Submit**.

Step 15 Configure the physical interface members as described in the "Assigning Physical Interfaces to a Port Channel" section on page 6-12.

After you create the port-channel interface, you need to assign physical interfaces to the port channel.

Assigning Physical Interfaces to a Port Channel

After you have configured a logical port-channel interface, you must assign multiple physical interfaces to the port channel. You can assign up to four physical interfaces to one port-channel interface, depending on the WAAS device.

You can assign an interface to a port channel only if the interface does not have an IP address assigned. The interface uses the IP address of the port channel.

You cannot combine built-in Ethernet ports with ports on a Cisco Interface Module into the same port-channel interface.

Note Removing a physical interface from a port channel on device models WAE-612/674/7341/7371 and WAVE-574 can cause network disruption for up to 30 seconds. The best practice is to make such changes when traffic interception is disabled or at an off-peak time when traffic disruption is acceptable.

To add an interface to a port channel, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Network Interfaces**. The Network Interfaces window for the device appears.

Step 3 Click the **Edit** icon next to the physical interface that you want to assign to a port channel. The Modifying Network Interface window appears.

Choose a physical interface, not a logical interface (standby, port channel, or BVI), in this step.

Step 4 Complete the following steps to assign the interface to a port channel:

a. In the Port Type To Assign drop-down list, choose **PortChannel**.

b. In the Port Channel Number drop-down list, choose the number of the port channel to which you want to add the physical interface.

Step 5 Click **Submit**.

Configuring a Load-Balancing Method for Port-Channel Interfaces

Before you configure load balancing, ensure that you have configured the port-channel settings described in the "Configuring Port-Channel Settings" section on page 6-9.

To configure load balancing, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Network > Port Channel**.

Step 3 From the Load Balancing Method drop-down list, choose a load-balancing method:

- **src-dst-ip-port**—The distribution function is based on a combination of source and destination IP addresses and ports. This load-balancing method is available only on devices running version 4.4.1 and later.
- **src-dst-ip**—The distribution function is based on a combination of source and destination IP addresses. This load-balancing method is available only on devices running version 5.0.1 and later.
- **round-robin**—Round robin allows traffic to be distributed evenly among all interfaces in the channel group. This load-balancing method is available only on devices running versions earlier than 4.4.1.

Step 4 Click **Submit**.

To configure a load-balancing method from the CLI, you can use the **port-channel** global configuration command.

Note A device group may be configured with a load-balancing method supported only by previous WAAS software versions to configure devices running previous versions. When viewing the Port Channel Settings page for a version 4.4.1 or later device that gets its settings from such a device group, you may see an unsupported load-balancing method listed. However, a version 4.4.1 or later device supports only the load-balancing methods as described above, regardless of what the device group or device configuration window shows for the setting.

Configuring Interfaces for DHCP

Note You must disable autoregistration before you can manually configure an interface for DHCP.

A WAAS device sends its configured client identifier and hostname to the DHCP server when requesting network information. You can configure DHCP servers to identify the client identifier information and the hostname information that the WAAS device is sending and then to send back the specific network settings that are assigned to the WAAS device.

To enable an interface for DHCP, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Network Interfaces**. The Network Interfaces listing window appears.

Step 3 Choose the physical interface that you want to modify and click the **Edit** taskbar icon. (For devices using WAAS versions earlier than 5.0, click the **Edit** icon next to the interface.)

The Interface Settings window appears.

Note Do not choose a logical interface (standby, port channel, or BVI) in this step, because you cannot configure DHCP on a logical interface. In addition, do not choose the internal interface (GigabitEthernet 1/0) on an NME-WAE or SM-SRE module, because this interface can be configured only through the host router CLI. For NME-WAE details, see the document *Configuring Cisco WAAS Network Modules for Cisco Access Routers*. For SM-SRE details, see the document *Cisco SRE Service Module Configuration and Installation Guide*.

Step 4 Check the **Use DHCP** check box.

When this check box is checked, the IP address and netmask fields are disabled.

Step 5 In the Hostname field, specify the hostname for the WAAS device or other device.

Step 6 In the Client Id field, specify the configured client identifier for the device.

The DHCP server uses this identifier when the WAAS device requests the network information for the device.

Step 7 Click **Submit**.

Modifying Virtual Interface Settings for a vWAAS Device

To modify the settings of an existing vWAAS interface, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Network Interfaces**.

The Network Interfaces window appears, listing the network interfaces configured.

Note Certain values (including autosense) are not applicable to a vWAAS interface.

Step 3 Choose the interface that you want to modify and click the **Edit** taskbar icon. (For devices using WAAS versions earlier than 5.0, click the **Edit** icon next to the interface.)

The Interface Settings window appears, displaying the interface configurations on a particular slot and port.

Note Interface configurations for slot, port, and port type are set for virtual interfaces during initial startup or by using the WAAS CLI.

Some of the fields in the window (port-channel number, autosense, speed, mode, and standby-related fields) are not available because they are not applicable.

Step 4 (Optional) In the Description field, optionally enter a description for the interface.

Step 5 (Optional) Check the **Use CDP** check box to enable the Cisco Discovery Protocol (CDP) on an interface.

When enabled, CDP obtains protocol addresses of neighboring devices and discovers the platform of those devices. It also shows information about the interfaces used by your router.

Configuring CDP from the CDP Settings window enables CDP globally on all the interfaces. For information on configuring CDP settings, see the "Configuring CDP Settings" section on page 6-25.

Step 6 (Optional) Check the **Shutdown** check box to shut down the virtual interface.

Step 7 (Optional) From the Load Interval drop-down list, choose the interval in seconds at which to poll the interface for statistics and calculate throughput. The default is 30 seconds. (The Load Interval item is not shown for devices using WAAS versions earlier than 5.0.)

Step 8 Specify a value (in bytes) in the MTU field to set the interface Maximum Transmission Unit (MTU) size.

The range is 576–1500 bytes. The MTU is the largest size of IP datagram that can be transferred using a specific data link connection.

Note The MTU field is not editable if a system jumbo MTU is configured.

Step 9 Check the **Use DHCP** check box to obtain an interface IP address through DHCP. Checking this box hides the IP address and Netmask fields. (For devices with WAAS versions earlier than 5.0, these fields are not hidden but become grayed out.)

a. (Optional) In the Hostname field, specify the hostname for the WAAS device or other device.

b. (Optional) In the Client Id field, specify the configured client identifier for the device. The DHCP server uses this identifier when the WAAS device requests the network information for the device.

Step 10 In the Address field, enter a new IP address to change the interface IP address.

Step 11 In the Netmask field, enter a new netmask to change the interface netmask.

Step 12 In the Default Gateway field, enter the default gateway IP address. The gateway interface IP address should be in the same network as one of the device's network interfaces. If an interface is configured for DHCP, then this field is read only. (The Default Gateway field is not shown for devices using WAAS versions 5.0 or later; instead, configure it as described in the "Configuring the Default Gateway" section on page 6-9.)

Step 13 (Optional) From the Inbound ACL drop-down list, choose an IP ACL to apply to inbound packets.

The drop-down list contains all the IP ACLs that you configured in the system.

Step 14 (Optional) From the Outbound ACL drop-down list, choose an IP ACL to apply to outbound packets.

Step 15 Click **OK**. (For devices using WAAS versions earlier than 5.0, click **Submit**.)

Configuring Optimization on WAAS Express Interfaces

WAAS Express device interfaces are configured by using the router CLI, not through the WAAS Central Manager. However, you can enable or disable WAAS optimization on the available interfaces on the router.

To enable or disable WAAS optimization on WAAS Express device interfaces, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *WAAS-Express-device-name* (or **Device Groups >** *WAAS-Express-device-group-name*).

Step 2 Choose **Configure > Network > Network Interfaces**. The Network Interfaces window appears and lists the available interfaces. (See Figure 6-2.)

Note Loopback interfaces are not included because they are not valid interfaces for optimization. Null, Virtual-Access, NVI, and Embedded-Service interfaces are also not supported.

Figure 6-2 WAAS Express Network Interfaces Device Window

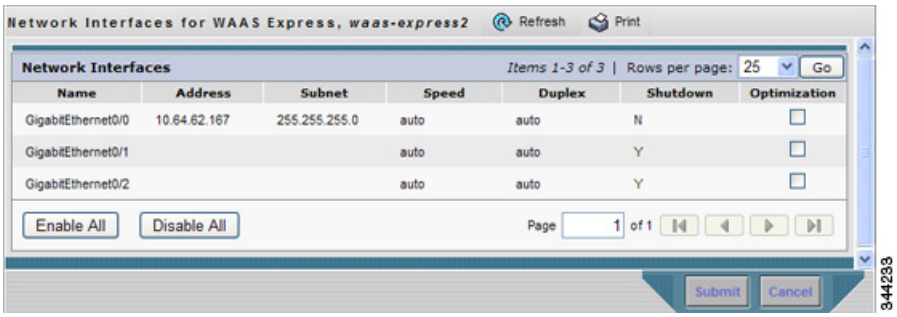

For a device group, the Network Interfaces window appears differently and displays an interface name, the number of devices that contain that interface, and the Optimization check box, which is checked if any devices in the group have optimization enabled on the interface. A message describes how many devices have optimization enabled on the interface. (See Figure 6-3.)

Figure 6-3 WAAS Express Network Interfaces Device Group Interfaces Window

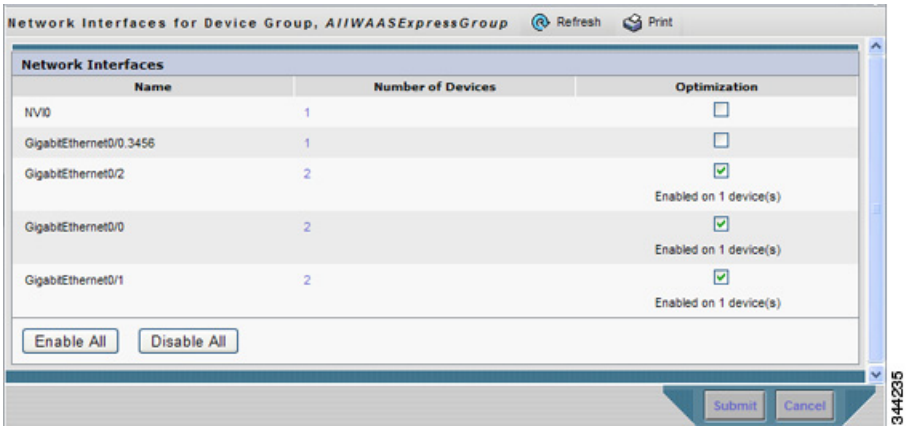

Step 3 Check the **Optimization** check box for each interface on which you want to enable WAAS optimization. Remove check marks from interfaces on which you want to disable WAAS optimization. You can click **Enable All** to select all interfaces or click **Disable All** to deselect all interfaces.

Enable WAAS optimization only on WAN interfaces, not LAN interfaces.

For a device group, checking the optimization check box for an interface enables optimization on that interface for all devices in the group that have the interface. You can click the number of devices to display a list of devices on which an interface is available and individually configure optimization on those devices. (See Figure 6-4.)

Figure 6-4 WAAS Express Network Interfaces Device Group Devices Window

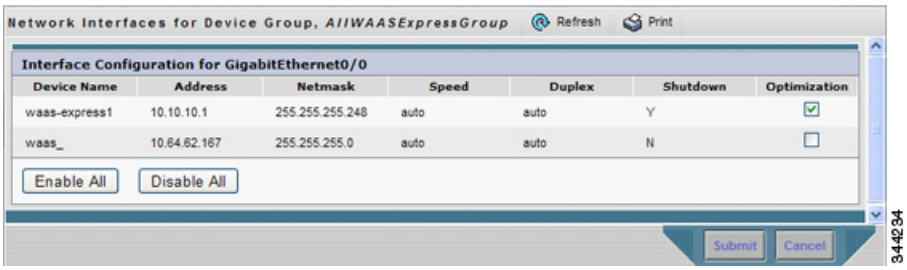

Step 4 Click **Submit**.

Bridging to a Virtual Blade Interface

To provide network connectivity to a virtual blade, you use a bridge group and bridge virtual interface (BVI) to associate a physical interface with a virtual interface on the virtual blade.

BVIs are supported only on WAAS devices that support virtual blades. BVIs are not supported on AppNav Controller Interface Modules or on WAAS devices operating as AppNav Controllers.

You can create up to five bridge interfaces on a device, depending on the device model.

Configuring a BVI differs, depending on the version of the WAAS device that you are configuring. See one of the following topics:

- Configuring a Bridge Virtual Interface on a Device with Version 5.0 or Later, page 6-17
- Configuring a Bridge Virtual Interface on a Device Earlier than Version 5.0, page 6-18

Configuring a Bridge Virtual Interface on a Device with Version 5.0 or Later

To configure a BVI for devices with WAAS version 5.0 or later, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Network Interfaces**. The Network Interfaces window for the device appears.

Step 3 In the lower part of the window, click the **Bridge** tab.

Step 4 In the taskbar of the lower area, click the **Create Bridge** icon. The Create Bridge window appears.

Step 5 From the Bridge Index drop-down list, choose a bridge group number for the interface.

Step 6 From the Protocol drop-down list, choose the **ieee** protocol type to support a BVI.

Step 7 (Optional) In the Description field, enter a description for the interface.

Step 8 (Optional) From the Load Interval drop-down list, choose the interval in seconds at which to poll the interface for statistics and calculate throughput. The default is 30 seconds.

Step 9 (Optional) Check the **Use DHCP** check box to obtain an interface IP address through DHCP. Checking this box hides the Address and Netmask fields.

Optionally supply a hostname in the Hostname field and a client ID in the Client Id field.

Step 10 In the Address field, specify the IP address of the interface.

Step 11 In the Netmask field, specify the netmask of the interface.

Step 12 (Optional) In the Secondary Address and Secondary Netmask fields, enter up to four secondary IP addresses and corresponding subnet masks.

Step 13 In the Assign Interfaces area, check the box next to the interface that you want to assign to this bridge group and click the **Assign** taskbar icon. To unassign an assigned interface, check the interface that you want to unassign and click the **Unassign** taskbar icon. Only one interface can be assigned to the bridge group and it can be a physical, port-channel, or standby interface.

Step 14 Click **OK**.

Configuring a Bridge Virtual Interface on a Device Earlier than Version 5.0

To configure a BVI for devices with WAAS versions earlier than 5.0, follow these steps:

1. Create a bridge group.
2. Create a bridge virtual interface in the bridge group.
3. Assign one physical, port-channel, or standby interface to the bridge group.
4. Assign the virtual blade interface to the bridge group.

These steps are described in more detail in this section.

To create a bridge group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Bridge**.

The Bridge Settings window appears, listing the bridge interfaces configured.

From the Bridge Settings window, you can perform the following tasks:

- Delete an existing bridge interface by clicking the **Edit** icon next to the interface number. You can then delete the bridge interface by clicking the **Delete** taskbar icon.
- Add a new bridge interface, as described in the following steps.

Step 3 Click the **Create Bridge Interface** taskbar icon to create a bridge interface.

The Creating new Bridge window appears.

Step 4 From the Bridge Index drop-down list, choose the number of the bridge interface (1–4).

Step 5 From the Protocol drop-down list, choose the **ieee** protocol type to support a BVI.

Step 6 Click **Submit**.

To create a bridge group from the CLI, you can use the **bridge** global configuration command.

After you create the bridge group, you must create a bridge virtual interface associated with the bridge group.

To create the bridge virtual interface, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Network Interfaces**. The Network Interfaces window appears, listing all the interfaces for the chosen device.

Step 3 In the taskbar, click the **Create New Interface** icon. The Creating New Network Interface window appears.

Step 4 From the Port Type drop-down list, choose **BVI**.

The window refreshes and provides fields for configuring the network interface settings.

Step 5 In the Bridge Group Number drop-down list, choose the number of the bridge group for this interface. Up to five bridge groups are supported, depending on the WAAS device model.

Step 6 (Optional) In the Description field, enter a description for the bridge virtual interface.

Step 7 Check the **Use DHCP** check box to obtain an interface IP address through DHCP. Checking this box grays out the IP address and Netmask fields.

a. (Optional) In the Hostname field, specify the hostname for the WAAS device or other device.

b. (Optional) In the Client Id field, specify the configured client identifier for the device. The DHCP server uses this identifier when the WAAS device requests the network information for the device.

Step 8 In the Default Gateway field, enter the default gateway IP address. If an interface is configured for DHCP, then this field is read only.

Step 9 In the Address field, specify the IP address of the interface.

Step 10 In the Netmask field, specify the netmask of the interface.

Step 11 In the Secondary Address and Secondary Netmask fields 1 through 4, enter up to four different IP addresses and secondary netmasks for the interface.

Step 12 Click **Submit**.

To create a bridge virtual interface from the CLI, you can use the **interface bvi** global configuration command.

After you create the bridge virtual interface, you must assign a physical, port-channel, or standby interface to the bridge group.

To assign an interface to the bridge group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Network Interfaces**. The Network Interfaces window appears, listing all the interfaces for the chosen device.

Step 3 Click the **Edit** icon next to the physical, port-channel, or standby interface that you want to assign to the bridge group.

Do not choose a primary interface because a primary interface cannot be assigned to a bridge group.

Step 4 In the Description field, optionally enter a description for the interface.

Step 5 Leave the Address and Netmask fields empty.

Step 6 If the interface is a physical interface, in the Port Type To Assign drop-down list, choose **Bridge Group**.

Step 7 In the Bridge Group Number drop-down list, choose the bridge group to which to assign the interface.

Step 8 Click **Submit**.

To assign a physical, port-channel, or standby interface to the bridge group from the CLI, you can use the **interface GigabitEthernet**, **interface TenGigabitEthernet**, **interface portchannel**, or **interface standby** global configuration commands, with the **bridge-group** keyword.

After you assign a physical or port-channel interface to the bridge group, you must assign a virtual blade interface to the bridge group. For details, see the "Configuring Virtual Blades" section on page 14-4.

Configuring Management Interface Settings

On WAAS devices with version 5.0 or later, you can designate a specific interface to be used as the management interface for communicating with the Central Manager, Telnet, SSH, and so on. This configuration separates management traffic from data traffic. If you designate a management interface, you must have another active interface to handle data traffic.

To configure the management interface settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Management Interface Settings**.

The Management Interface Settings window appears.

Step 3 From the Management Interface drop-down list, choose the interface that you want to use as the management interface.

Step 4 In the Management Default Gateway field, enter the default gateway IP address for management traffic.

Step 5 Check the **Use Management Interface for FTP Traffic** check box if you want to use the designated management interface for FTP traffic.

Step 6 Check the **Use Management Interface for TFTP Traffic** check box if you want to use the designated management interface for TFTP traffic.

Step 7 Click **Submit**. A confirmation message appears.

Step 8 Click **OK**.

To configure a different default gateway for management traffic from the CLI, you can use the **ip default-gateway management** global configuration command.

When you have designated a management interface, you can create static IP routes for management traffic, so that any IP packet that is designated for the specified destination uses the configured route.

To configure a static route for management traffic, follow these steps:

Step 1 In the Management Interface Settings window, in the Management IP Routes area, click the **Create Management IP Route** taskbar button. The Management IP Routes window appears.

Step 2 In the Destination Network Address field, enter the destination network IP address.

Step 3 In the Netmask field, enter the destination host netmask.

Step 4 In the Gateway's IP Address field, enter the IP address of the gateway interface.

The gateway interface IP address should be in the same network as the device's management interface.

Step 5 Click **Submit**.

To configure a static route for management traffic from the CLI, you can use the **ip route management** global configuration command.

Configuring a Jumbo MTU

A jumbo MTU can be configured on the following devices: WAE-674/7341/7371, WAVE-294/594/694/7541/7571/8541, and vWAAS.

If configured, a jumbo MTU applies to all the device interfaces, including logical interfaces with at least one member physical interface. The MTU for individual interfaces cannot be changed while the jumbo MTU is configured. If the jumbo MTU is disabled, all interfaces are configured with a MTU of 1500.

To configure a jumbo MTU, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Jumbo MTU**.

The Jumbo MTU Settings window appears.

Step 3 In the System Jumbo MTU field, enter the jumbo MTU size in bytes (maximum size varies by platform).

Step 4 Click **Submit**.

Note If the original and optimized maximum segment sizes are set to their default values and you configure a jumbo MTU setting, the segment sizes are changed to the jumbo MTU setting minus 68 bytes. If you have configured custom maximum segment sizes, their values are not changed if you configure a jumbo MTU. For more information on configuring maximum segment sizes, see the "Modifying the Acceleration TCP Settings" section on page 13-61.

To configure a jumbo MTU from the CLI, you can use the **system jumbomtu** global configuration command.

Configuring TCP Settings

For data transactions and queries between client and servers, the size of windows and buffers is important, so fine-tuning the TCP stack parameters becomes the key to maximizing cache performance.

Because of the complexities involved in TCP parameters, be careful when tuning these parameters. In nearly all environments, the default TCP settings are adequate. Fine-tuning TCP settings is for network administrators with adequate experience and full understanding of TCP operation details.

To configure TCP and IP settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Network > TCP/IP Settings > TCP/IP**. The TCP/IP Settings window appears.

Step 3 Make the necessary changes to the TCP settings.

See Table 6-1 for a description of each TCP field in this window.

Step 4 Click **Submit**.

A "Click Submit to Save" message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking **Reset**. The Reset button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

Table 6-1 TCP Settings

TCP Setting	Description
TCP General Settings	
Enable Explicit Congestion Notification	Enables reduction of delay and packet loss in data transmissions. It provides TCP support for RFC 2581. By default, this option is enabled. For more information, see the "Explicit Congestion Notification" section on page 6-23.
Initial Send Congestion Window Size	Initial congestion window size value in segments. The range is 0 to 10 segments. The default is 0 segments. For more information, see the "Congestion Windows" section on page 6-23.
ReTransmit Time Multiplier	Factor used to modify the length of the retransmit timer by 1 to 3 times the base value determined by the TCP algorithm. The default is 1, which leaves the times unchanged. The range is 1 to 3. For more information, see the "Retransmit Time Multiplier" section on page 6-23. **Note** Modify this factor with caution. It can improve throughput when TCP is used over slow reliable connections but should never be changed in an unreliable packet delivery environment.
Keepalive Probe Count	Number of times that the WAAS device can retry a connection before the connection is considered unsuccessful. The range is 1 to 120 attempts. The default is 4 attempts.
Keepalive Probe Interval	Length of time that the WAAS device keeps an idle connection open. The default is 75 seconds.
Keepalive Timeout	Length of time that the WAAS device keeps a connection open before disconnecting. The range is 1 to 120 seconds. The default is 90 seconds.
Enable Path MTU Discovery	Enables discovery of the largest IP packet size allowable between the various links along the forwarding path and automatically sets the correct value for the packet size. By default, this option is disabled. For more information, see the "Path MTU Discovery" section on page 6-24.

To configure TCP settings from the CLI, you can use the **tcp** global configuration command.

To enable the MTU discovery utility from the CLI, you can use the **ip path-mtu-discovery enable** global configuration command.

This section contains the following topics:

Explicit Congestion Notification

The TCP Explicit Congestion Notification (ECN) feature allows an intermediate router to notify the end hosts of impending network congestion. It also provides enhanced support for TCP sessions associated with applications that are sensitive to delay or packet loss. The major issue with ECN is that the operation of both the routers and the TCP software stacks needs to be changed to accommodate the operation of ECN.

Congestion Windows

The congestion window (*cwnd*) is a TCP state variable that limits the amount of data that a TCP sender can transmit onto the network before receiving an acknowledgment (ACK) from the receiving side of the TCP transmission. The TCP *cwnd* variable is implemented by the TCP congestion avoidance algorithm. The goal of the congestion avoidance algorithm is to continually modify the sending rate so that the sender automatically senses any increase or decrease in available network capacity during the entire data flow. When congestion occurs (manifested as packet loss), the sending rate is first lowered then gradually increased as the sender continues to probe the network for additional capacity.

Retransmit Time Multiplier

The TCP sender uses a timer to measure the time that has elapsed between sending a data segment and receiving the corresponding ACK from the receiving side of the TCP transmission. When this retransmit timer expires, the sender (according to the RFC standards for TCP congestion control) must reduce its sending rate. However, because the sender is not reducing its sending rate in response to network congestion, the sender is not able to make any valid assumptions about the current state of the network. Therefore, in order to avoid congesting the network with an inappropriately large burst of data, the sender implements the slow start algorithm, which reduces the sending rate to one segment per transmission. (See the "TCP Slow Start" section on page 6-24.)

You can modify the sender's retransmit timer by using the Retransmit Time Multiplier field in the WAAS Central Manager. The retransmit time multiplier modifies the length of the retransmit timer by one to three times the base value, as determined by the TCP algorithm that is being used for congestion control.

When making adjustments to the retransmit timer, be aware that they affect performance and efficiency. If the retransmit timer is triggered too early, the sender pushes duplicate data onto the network unnecessarily; if the timer is triggered too slowly, the sender remains idle for too long, unnecessarily slowing data flow.

TCP Slow Start

Slow start is one of four congestion control algorithms used by TCP. The slow start algorithm controls the amount of data being inserted into the network at the beginning of a TCP session when the capacity of the network is not known.

For example, if a TCP session began by inserting a large amount of data into the network, much of the initial burst of data would likely be lost. Instead, TCP initially transmits a modest amount of data that has a high probability of successful transmission. Next, TCP probes the network by sending increasing amounts of data as long as the network does not show signs of congestion.

The slow start algorithm begins by sending packets at a rate that is determined by the congestion window, or *cwnd* variable. (See the "Congestion Windows" section on page 6-23.) The algorithm continues to increase the sending rate until it reaches the limit set by the slow start threshold (*ssthresh*) variable. Initially, the value of the *ssthresh* variable is adjusted to the receiver's maximum segment size (RMSS). However, when congestion occurs, the *ssthresh* variable is set to half the current value of the *cwnd* variable, marking the point of the onset of network congestion for future reference.

The starting value of the *cwnd* variable is set to that of the sender maximum segment size (SMSS), which is the size of the largest segment that the sender can transmit. The sender sends a single data segment, and because the congestion window is equal to the size of one segment, the congestion window is now full. The sender then waits for the corresponding ACK from the receiving side of the transmission. When the ACK is received, the sender increases its congestion window size by increasing the value of the *cwnd* variable by the value of one SMSS. Now the sender can transmit two segments before the congestion window is again full and the sender is once more required to wait for the corresponding ACKs for these segments. The slow start algorithm continues to increase the value of the *cwnd* variable and therefore increase the size of the congestion window by one SMSS for every ACK received. If the value of the *cwnd* variable increases beyond the value of the *ssthresh* variable, then the TCP flow control algorithm changes from the slow start algorithm to the congestion avoidance algorithm.

Path MTU Discovery

The WAAS software supports the IP Path Maximum Transmission Unit (MTU) Discovery method, as defined in RFC 1191. When enabled, the Path MTU Discovery feature discovers the largest IP packet size allowable between the various links along the forwarding path and automatically sets the correct value for the packet size. By using the largest MTU that the links can handle, the sending device can minimize the number of packets it must send.

IP Path MTU Discovery is useful when a link in a network goes down, which forces the use of another, different MTU-sized link. IP Path MTU Discovery is also useful when a connection is first being established, and the sender has no information about the intervening links.

Note IP Path MTU Discovery is a process initiated by the sending device. If a server does not support IP Path MTU Discovery, the receiving device will have no available means to avoid fragmenting datagrams generated by the server.

By default, this feature is disabled. With the feature disabled, the sending device uses a packet size that is the lesser of 576 bytes and the next hop MTU. Existing connections are not affected when this feature is turned on or off.

Configuring Static IP Routes

The WAAS software allows you to configure a static route for a network or host. Any IP packet designated for the specified destination uses the configured route.

To configure a static IP route, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Network > TCP/IP Settings > Static Routes**. The IP Route Entries window appears.

Step 3 In the taskbar, click the **Create New IP Route Entry** icon. The Creating New IP Route window appears.

Step 4 In the Destination Network Address field, enter the destination network IP address.

Step 5 In the Netmask field, enter the destination host netmask.

Step 6 In the Gateway's IP Address field, enter the IP address of the gateway interface.

The gateway interface IP address should be in the same network as that of one of the device's network interfaces.

Step 7 Click **Submit**.

To configure a static route from the CLI, you can use the **ip route** global configuration command.

Aggregating IP Routes

An individual WAE device can have IP routes defined and can belong to device groups that have other IP routes defined.

In the IP Route Entries window, the Aggregate Settings radio button controls how IP routes are aggregated for an individual device, as follows:

- Choose **Yes** if you want to configure the device with all IP routes that are defined for itself and for device groups to which it belongs.
- Choose **No** if you want to limit the device to just the IP routes that are defined for itself.

When you change the setting, you get the following confirmation message: "This option will take effect immediately and will affect the device configuration. Do you wish to continue?" Click **OK** to continue.

Configuring CDP Settings

The Cisco Discovery Protocol (CDP) is a device discovery protocol that runs on all Cisco-manufactured devices. With CDP, each device in a network sends periodic messages to all other devices in the network. All devices listen to periodic messages that are sent by others to learn about neighboring devices and determine the status of their interfaces.

With CDP, network management applications can learn the device type and the Simple Network Management Protocol (SNMP) agent address of neighboring devices. Applications are able to send SNMP queries within the network. CiscoWorks2000 also discovers the WAAS devices by using the CDP packets that are sent by the WAAS device after booting.

To perform device-related tasks, the WAAS device platform must support CDP to be able to notify the system manager of the existence, type, and version of the WAAS device platform.

To configure CDP settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Network > CDP**. The CDP Settings window appears.

Step 3 Check the **Enable** check box to enable CDP support. By default, this option is enabled.

Step 4 In the Hold Time field, enter the time (in seconds) to specify the length of time that a receiver is to keep the CDP packets.

The range is 10 to 255 seconds. The default is 180 seconds.

Step 5 In the Packet Send Rate field, enter a value (in seconds) for the interval between CDP advertisements.

The range is 5 to 254 seconds. The default is 60 seconds.

Step 6 Click **Submit**.

To configure CDP settings from the CLI, you can use the **cdp** global configuration command.

Configuring the DNS Server

DNS allows the network to translate domain names entered in requests into their associated IP addresses. To configure DNS on a WAAS device, you must complete the following tasks:

- Specify the list of DNS servers, which are used by the network to translate requested domain names into IP addresses that the WAAS device should use for domain name resolution.
- Enable DNS on the WAAS device.

To configure DNS server settings for a WAAS device, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Network > DNS**. The DNS Settings window appears.

Step 3 In the Local Domain Name field, enter the name of the local domain. You can configure up to three local domain names. Separate items in the list with a space.

Step 4 In the List of DNS Servers field, enter a list of DNS servers used by the network to resolve hostnames to IP addresses.

You can configure up to three DNS servers. Separate items in the list with a space.

Step 5 Click **Submit**.

A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default and device group settings. To revert to the previously configured window settings, click **Reset**. The Reset button appears only when you have applied default or group settings to change the current device settings but the settings have not yet been submitted.

To configure DNS name servers from the CLI, you can use the **ip name-server** global configuration command.

Configuring Windows Name Services

To configure Windows name services for a device or device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Network > WINS**. The Windows Name Services Settings window appears.

Step 3 In the Workgroup or Domain Name field, enter the name of the workgroup (or domain) in which the chosen device or device group resides.

This name must be entered in shortname format and cannot exceed 127 characters. Valid characters include alphanumeric characters, a forward slash (\), an underscore (_), and a dash (-).

For example, if your domain name is cisco.com, the short name format is cisco.

Step 4 Check the **NT** check box if the workgroup or domain is a Windows NT 4 domain. For example, if your domain name is cisco.com, the short name format is cisco. If your workgroup or domain is a Windows 2000 or Windows 2003 domain, do not check the NT check box. By default, this option is disabled.

Step 5 In the WINS server field, enter the hostname or IP address of the Windows Internet Naming Service (WINS) server.

Step 6 Click **Submit**.

To configure Windows name services from the CLI, you can use the **windows-domain** global configuration command.

Configuring Directed Mode

By default, WAAS transparently sets up new TCP connections to peer WAEs, which can cause firewall traversal issues when a WAAS device tries to optimize the traffic. If a WAE device is behind a firewall that prevents traffic optimization, you can use the directed mode of communicating to a peer WAE. In directed mode, all TCP traffic that is sent to a peer WAE is encapsulated in UDP, which allows a firewall to either bypass the traffic or inspect the traffic (by adding a UDP inspection rule).

Any firewall between two WAE peers must be configured to pass UDP traffic on port 4050, or whatever custom port is configured for directed mode if a port other than the default is used. Additionally, because the WAAS automatic discovery process uses TCP options before directed mode begins sending UDP traffic, the firewall must be configured to pass the TCP options. Cisco firewalls can be configured to allow TCP options by using the **ip inspect waas** command (for Cisco IOS Release 12.4(11)T2 and later releases) or the **inspect waas** command (for FWSM 3.2(1) and later releases and PIX 7.2(3) and later releases).

After directed mode is activated, the WAE transparently intercepts only packets coming from the LAN, while WAN packets are directly routed between the WAEs using UDP.

Directed mode operates with all configurable methods of traffic interception. Directed mode requires that you configure the WAAS devices (or inline interfaces) with routable, non-NATed IP addresses. When using directed mode with inline mode, you must configure the inline group with routable IP addresses on its interfaces or traffic is black holed.

If a WAE at either end of a peer WAE connection specifies directed mode, and both WAEs support directed mode, then both WAEs use directed mode, even if one is not explicitly configured for directed mode. If a peer WAE does not support directed mode, then the peers pass through traffic unoptimized and each WAE creates a transaction log entry that notes the failed directed mode attempt.

You can invoke directed mode operation in the following ways:

- Directed mode can be explicitly activated in the WAAS Central Manager or by CLI.
- Directed mode can be automatically invoked when a peer WAE requests that directed mode be used.

To activate directed mode, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Network > Directed Mode**. The Directed Mode Settings window appears.

Step 3 Check the **Enable directed mode** check box to activate directed mode.

Step 4 In the UDP Port field, enter a port number to configure a custom UDP port for directed mode. The default is port 4050.

Step 5 Click **Submit** to save the settings.

To configure directed mode from the CLI, use the **directed-mode** global configuration command.

CHAPTER 7

Configuring Administrative Login Authentication, Authorization, and Accounting

This chapter describes how to configure administrative login authentication, authorization, and accounting for Wide Area Application Services (WAAS) devices.

This chapter contains the following sections:

You use the WAAS Central Manager GUI to centrally create and manage two different types of administrator user accounts (device-based CLI accounts and roles-based accounts) for your WAAS devices. For more information, see Chapter 8, "Creating and Managing Administrator User Accounts and Groups."

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

About Administrative Login Authentication and Authorization

In the WAAS network, administrative login authentication and authorization are used to control login requests from administrators who want to access a WAAS device for configuring, monitoring, or troubleshooting purposes.

Login authentication is the process by which WAAS devices verify whether the administrator who is attempting to log in to the device has a valid username and password. The administrator who is logging in must have a user account registered with the device. User account information serves to authorize the user for administrative login and configuration privileges. The user account information is stored in an AAA database, and the WAAS devices must be configured to access the particular authentication server (or servers) where the AAA database is located. When the user attempts to login to a device, the device compares the person's username, password, and privilege level to the user account information that is stored in the database.

The WAAS software provides the following authentication, authorization, and accounting (AAA) support for users who have external access servers (for example, RADIUS or TACACS+ servers), and for users who need a local access database with AAA features:

- Authentication (or login authentication) is the action of determining who the user is. It checks the username and password.
- Authorization (or configuration) is the action of determining what a user is allowed to do. It permits or denies privileges for authenticated users in the network. Generally, authentication precedes authorization. Both authentication and authorization are required for a user log in.
- Accounting is the action of keeping track of administrative user activities for system accounting purposes. In the WAAS software, AAA accounting through TACACS+ is supported. For more information, see the "Configuring AAA Accounting for WAAS Devices" section on page 7-31.

Note An administrator can log in to the WAAS Central Manager device through the console port or the WAAS Central Manager GUI. An administrator can log in to a WAAS device that is functioning as a data center or branch WAE through the console port or the WAE Device Manager GUI.

When the system administrator logs in to a WAAS device before authentication and authorization have been configured, the administrator can access the WAAS device by using the predefined superuser account (the predefined username is admin and the predefined password is default). When you log in to a WAAS device using this predefined superuser account, you are granted access to all the WAAS services and entities in the WAAS system.

Note Each WAAS device must have one administrator account with the username admin. You cannot change the username of the predefined superuser account. The predefined superuser account must have the username admin.

After you have initially configured your WAAS devices, we strongly recommend that you immediately change the password for the predefined superuser account (the predefined username is admin, the password is default, and the privilege level is superuser, privilege level 15) on each WAAS device.

For instructions on using the WAAS Central Manager GUI to change the password for the predefined superuser account, see the "Changing the Password for Your Own Account" section on page 8-6.

Figure 7-1 shows how an administrator can log in to a WAE through the console port or the WAAS GUIs (the WAAS Central Manager GUI or the WAE Device Manager GUI). When the WAAS device receives an administrative login request, the WAE can check its local database or a remote third-party database (TACACS+, RADIUS, or Windows domain database) to verify the username with the password and to determine the access privileges of the administrator.

Figure 7-1 Authentication Databases and a WAE

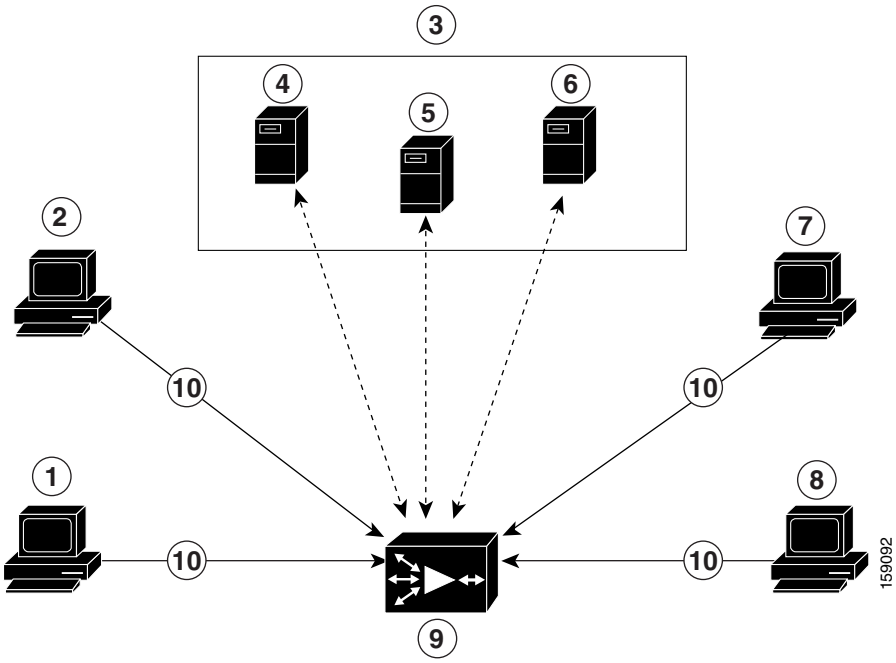

1	FTP/SFTP client	**6**	Windows domain server
2	WAAS Central Manager GUI or WAE Device Manager GUI	**7**	Console or Telnet clients
3	Third-party AAA servers	**8**	SSH client
4	RADIUS server	**9**	WAE that contains a local database and the default primary authentication database
5	TACACS+ server	**10**	Administrative login requests

The user account information is stored in an AAA database, and the WAAS devices must be configured to access the particular authentication server (or servers) that contains the AAA database. You can configure any combination of these authentication and authorization methods to control administrative login access to a WAAS device:

- Local authentication and authorization
- RADIUS
- TACACS+
- Windows domain authentication

Note If you configure authentication using an external authentication server, you still must create a role-based user or user group account in the WAAS Central Manager as described in Chapter 8, "Creating and Managing Administrator User Accounts and Groups."

For more information on the default AAA configuration, see the "Default Administrative Login Authentication and Authorization Configuration" section on page 7-4. For more information on configuring AAA, see the "Configuring Administrative Login Authentication and Authorization" section on page 7-5.

Default Administrative Login Authentication and Authorization Configuration

By default, a WAAS device uses the local database to obtain login authentication and authorization privileges for administrative users.

Table 7-1 lists the default configuration for administrative login authentication and authorization.

Table 7-1 Default Configuration for Administrative Login Authentication and Authorization

Feature	Default Value
Administrative login authentication	Enabled
Administrative configuration authorization	Enabled
Authentication server failover because the authentication server is unreachable	Disabled
TACACS+ port	Port 49
TACACS+ login authentication (console and Telnet)	Disabled
TACACS+ login authorization (console and Telnet)	Disabled
TACACS+ key	None specified
TACACS+ server timeout	5 seconds
TACACS+ retransmit attempts	2 times
RADIUS login authentication (console and Telnet)	Disabled
RADIUS login authorization (console and Telnet)	Disabled
RADIUS server IP address	None specified
RADIUS server UDP authorization port	Port 1645
RADIUS key	None specified
RADIUS server timeout	5 seconds
RADIUS retransmit attempts	2 times
Windows domain login authentication	Disabled
Windows domain login authorization	Disabled
Windows domain password server	None specified
Windows domain realm (Kerberos realm used for authentication when Kerberos authentication is used). **Note** When Kerberos authentication is enabled, the default **realm** is DOMAIN.COM and security is the Active Directory Service (ADS).	Null string
Hostname or IP address of the Windows Internet Naming Service (WIN) server for Windows domain	None specified

Table 7-1 Default Configuration for Administrative Login Authentication and Authorization (continued)

Feature	Default Value
Window domain administrative group	There are no predefined administrative groups.
Windows domain NETBIOS name	None specified
Kerberos authentication	Disabled
Kerberos server hostname or IP address (host that is running the Key Distribution Center (KDC) for the given Kerberos realm	None specified
Kerberos server port number (port number on the KDC server)	Port 88
Kerberos local realm (default realm for WAAS)	kerberos-realm: null string
Kerberos realm (maps a hostname or DNS domain name to a Kerberos realm)	Null string

Note If you configure a RADIUS or TACACS+ key on the WAAS device (the RADIUS and the TACACS+ client), make sure that you configure an identical key on the external RADIUS or TACACS+ server.

You change these defaults through the WAAS Central Manager GUI, as described in the "Configuring Administrative Login Authentication and Authorization" section on page 7-5.

Multiple Windows domain utilities are included in the WAAS software to assist with Windows domain authentication configuration. You can access these utilities through the WAAS CLI by using the **windows-domain diagnostics** EXEC command.

Configuring Administrative Login Authentication and Authorization

To centrally configure administrative login authentication and authorization for a WAAS device or a device group (a group of WAEs), follow these steps:

Step 1 Determine the login authentication scheme that you want to configure the WAAS device to use when authenticating administrative login requests (for example, use the local database as the primary login database and your RADIUS server as the secondary authentication database).

Step 2 Configure the login access control settings for the WAAS device, as described in the "Configuring Login Access Control Settings for WAAS Devices" section on page 7-7.

Step 3 Configure the administrative login authentication server settings on the WAAS device (if a remote authentication database is to be used). For example, specify the IP address of the remote RADIUS servers, TACACS+ servers, or Windows domain server that the WAAS device should use to authenticate administrative login requests, as described in the following sections:

- Configuring RADIUS Server Authentication Settings, page 7-12
- About TACACS+ Server Authentication Settings, page 7-14

- Configuring Windows Domain Server Authentication Settings, page 7-17

Step 4 Specify one or all of the following login authentication configuration schemes that the WAAS device should use to process administrative login requests:

- Specify the administrative login authentication scheme.
- Specify the administrative login authorization scheme.
- Specify the failover scheme for the administrative login authentication server (optional).

For example, specify which authentication database the WAAS device should check to process an administrative login request. See the "Enabling Administrative Login Authentication and Authorization Schemes for WAAS Devices" section on page 7-26.

Caution Make sure that RADIUS, TACACS+, or Windows domain authentication is configured and operating correctly before disabling local authentication and authorization. If you disable local authentication and RADIUS, TACACS+, or Windows domain settings are not configured correctly, or if the RADIUS, TACACS+, or Windows domain server is not online, you may be unable to log in to the WAAS device.

You can enable or disable the local and the remote databases (TACACS+, RADIUS, and Windows domain) through the WAAS Central Manager GUI or the WAAS CLI. The WAAS device verifies whether all databases are disabled and, if so, sets the system to the default state (see Table 7-1). If you have configured the WAAS device to use one or more of the external third-party databases (TACACS+, RADIUS, or Windows domain authentication) for administrative authentication and authorization, make sure that you have also enabled the local authentication and authorization method on the WAAS device, and that the local method is specified as the last option; otherwise, the WAAS device will not go to the local authentication and authorization method by default if the specified external third-party databases are not reachable.

By default, local login authentication is enabled first. Local authentication and authorization uses locally configured login and passwords to authenticate administrative login attempts. The login and passwords are local to each WAAS device and are not mapped to individual usernames. When local authentication is disabled, if you disable all other authentication methods, local authentication is reenabled automatically.

You can disable local login authentication only after enabling one or more of the other administrative login authentication methods. However, when local login authentication is disabled, if you disable all other administrative login authentication methods, local login authentication is reenabled automatically. You cannot specify different administrative login authentication methods for console and Telnet connections.

We strongly recommend that you set the administrative login authentication and authorization methods in the same order. For example, configure the WAAS device to use RADIUS as the primary login method, TACACS+ as the secondary login method, Windows as the tertiary method, and the local method as the quaternary method for both administrative login authentication and authorization.

Note A TACACS+ server will not authorize a user who is authenticated by a different method. For example, if you configure Windows as the primary authentication method, but TACACS+ as the primary authorization method, TACACS+ authorization will fail.

We strongly recommend that you specify the local method as the last method in your prioritized list of login authentication and authorization methods. By adhering to this practice, if the specified external third-party servers (TACACS+, RADIUS, or Windows domain servers) are not reachable, a WAAS administrator can still log in to a WAAS device through the local authentication and authorization method.

This section describes how to centrally configure administrative login authentication and contains the following topics:

Configuring Login Access Control Settings for WAAS Devices

This section describes how to centrally configure remote login and access control settings for a WAAS device or device group and contains the following topics:

Configuring Secure Shell Settings for WAAS Devices

Secure Shell (SSH) consists of a server and a client program. Like Telnet, you can use the client program to remotely log in to a machine that is running the SSH server, but unlike Telnet, messages transported between the client and the server are encrypted. The functionality of SSH includes user authentication, message encryption, and message authentication.

Note By default, the SSH feature is disabled on a WAAS device.

The SSH management window in the WAAS Central Manager GUI allows you to specify the key length, login grace time, and maximum number of password guesses allowed when logging in to a specific WAAS device or device group for configuration, monitoring, or troubleshooting purposes.

To centrally enable the SSH feature on a WAAS device or a device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Network > Console Access > SSH**.

The SSH Configuration window appears. (See Figure 7-2.)

Figure 7-2 SSH Configuration Window

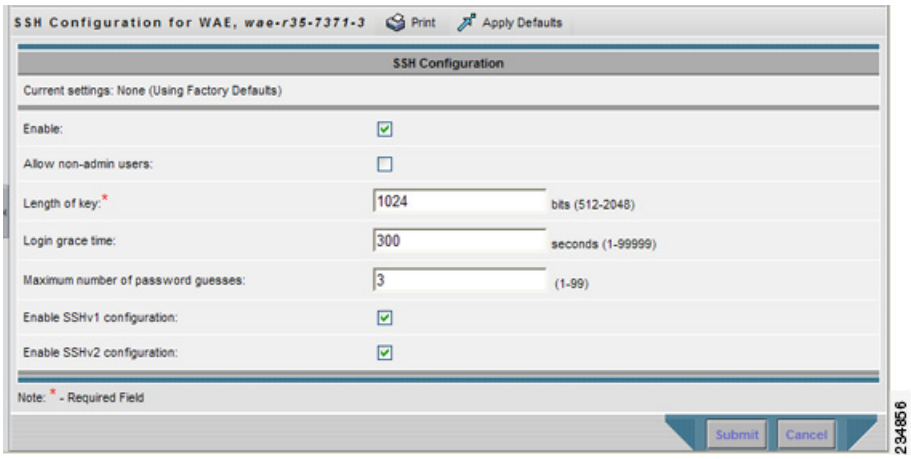

Step 3 Check the **Enable** check box to enable the SSH feature. SSH enables login access to the chosen WAAS device (or the device group) through a secure and encrypted channel.

Step 4 Check the **Allow non-admin users** check box to allow non-administrative users to gain SSH access to the chosen device (or device group). By default, this option is disabled.

Note Nonadministrative users are non-superuser administrators. All non-superuser administrators only have restricted access to a WAAS device because their login accounts have a privilege level of 0. Superuser administrators have full access to a WAAS device because their login accounts have the highest level of privileges, a privilege level of 15.

Step 5 In the Length of key field, specify the number of bits needed to create an SSH key. The default is 1024.

When you enable SSH, be sure to generate both a private and a public host key, which client programs use to verify the server's identity. When you use an SSH client and log in to a WAAS device, the public key for the SSH daemon running on the device is recorded in the client machine known_hosts file in your home directory. If the WAAS administrator subsequently regenerates the host key by specifying the number of bits in the Length of key field, you must delete the old public key entry associated with the WAAS device in the known_hosts file before running the SSH client program to log in to the WAAS device. When you use the SSH client program after deleting the old entry, the known_hosts file is updated with the new SSH public key for the WAAS device.

Step 6 In the Login grace time field, specify the number of seconds for which an SSH session will be active during the negotiation (authentication) phase between client and server before it times out. The default is 300 seconds.

Step 7 In the Maximum number of password guesses field, specify the maximum number of incorrect password guesses allowed per connection. The default is 3.

Although the value in the Maximum number of password guesses field specifies the number of allowed password guesses from the SSH server side, the actual number of password guesses for an SSH login session is determined by the combined number of allowed password guesses of the SSH server and the SSH client. Some SSH clients limit the maximum number of allowed password guesses to three (or to one in some cases), even though the SSH server allows more than this number of guesses. When you specify *n* allowed password guesses, certain SSH clients interpret this number as *n* + 1. For example, when configuring the number of guesses to two for a particular device, SSH sessions from some SSH clients will allow three password guesses.

Step 8 Specify whether the clients should be allowed to connect using the SSH protocol Version 1 or Version 2:

- To specify Version 1, check the **Enable SSHv1** check box.
- To specify Version 2, check the **Enable SSHv2** check box.

Note You can enable both SSH Version 1 and Version 2, or you can enable one version and not the other. You cannot disable both versions of SSH unless you disable the SSH feature by unchecking the **Enable** check box. (See Step 3.)

Step 9 Click **Submit** to save the settings.

A "Click Submit to Save" message appears in red in the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking the **Reset** button. The Reset button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

If you try to exit this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box only appears if you are using the Internet Explorer browser.

To configure SSH settings from the CLI, you can use the **sshd** and **ssh-key-generate** global configuration commands.

Disabling and Reenabling the Telnet Service for WAAS Devices

By default, the Telnet service is enabled on a WAAS device. You must use a console connection instead of a Telnet session to define device network settings on a WAAS device. However, after you have used a console connection to define the device network settings, you can use a Telnet session to perform subsequent configuration tasks.

You must enable the Telnet service before you can use the Telnet button in the Device Dashboard window to Telnet to a device.

Note Telnet is not supported in Internet Explorer. If you want to use the Telnet button from the Device Dashboard, use a different web browser.

To centrally disable the Telnet service on a WAAS device or a device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Network > Console Access > Telnet**. The Telnet Settings window appears.

Step 3 Uncheck the **Telnet Enable** check box to disable the terminal emulation protocol for remote terminal connection for the chosen device (or device group).

Step 4 Click **Submit** to save the settings.

A "Click Submit to Save" message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking the **Reset** button. The **Reset** button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

If you try to exit this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box only appears if you are using the Internet Explorer browser.

To centrally reenable the Telnet service on the device (or device group) at a later time, check the **Telnet Enable** check box in the Telnet Settings window and click **Submit**.

From the CLI, you can use the **no telnet enable** global configuration command to disable Telnet or the **telnet enable** global configuration command to enable it.

Configuring Message of the Day Settings for WAAS Devices

The Message of the Day (MOTD) feature enables you to provide information bits to the users when they log in to a device that is part of your WAAS network. There are three types of messages that you can set up:

- MOTD banner
- EXEC process creation banner
- Login banner

Note When you run an SSH version 1 client and log in to the device, the MOTD and login banners are not displayed. You need to use SSH version 2 to display the banners when you log in to the device.

To configure the MOTD settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > Console Access > Message of the day**. The MOTD Configuration window for the chosen device appears.

Step 3 To enable the MOTD settings, check the **Enable** check box. The Message of the Day (MOTD) banner, EXEC process creation banner, and Login banner fields become enabled.

Step 4 In the Message of the Day (MOTD) Banner field, enter a string that you want to display as the MOTD banner after a user logs in to the device.

Note In the Message of the Day (MOTD) Banner, EXEC Process Creation Banner, and Login Banner fields, you can enter a maximum of 1024 characters. A new line character (or Enter) is counted as two characters, as it is interpreted as \n by the system. You cannot use special characters such as `, % ,^ , and " in the MOTD text. If your text contains any of these special characters, WAAS software removes it from the MOTD output.

Step 5 In the EXEC Process Creation Banner field, enter a string to be displayed as the EXEC process creation banner when a user enters into the EXEC shell of the device.

Step 6 In the Login Banner field, enter a string to be displayed after the MOTD banner, when a user attempts to login to the device.

Step 7 To save the configuration, click **Submit**.

Configuring Exec Timeout Settings for WAAS Devices

To centrally configure the length of time that an inactive Telnet session remains open on a WAAS device or device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Click the **Edit** icon next to the device (or device group) for which you want to configure the EXEC timeout.

Step 3 Choose **Configure > Network > Console Access > Exec Timeout**.

Step 4 In the Exec Timeout field, specify the number of minutes after which an active session times out. The default is 15 minutes.

A Telnet session with a WAAS device can remain open and inactive for the period specified in this field. When the EXEC timeout period elapses, the WAAS device automatically closes the Telnet session.

Step 5 Click **Submit** to save the settings.

A "Click Submit to Save" message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking the **Reset** button. The **Reset** button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

If you try to exit this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box only appears if you are using the Internet Explorer browser.

To configure the Telnet session timeout from the CLI, you can use the **exec-timeout** global configuration command.

Configuring Line Console Carrier Detection for WAAS Devices

You need to enable carrier detection if you plan to connect the WAAS device to a modem for receiving calls.

Note By default, this feature is disabled on a WAAS device.

To centrally enable console line carrier detection for a WAAS device or device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Network > Console Access > Console Carrier Detect**. The Console Carrier Detect Settings window appears.

Step 3 Check the **Enable console line carrier detection before writing to the console** check box to enable the window for configuration.

Step 4 Click **Submit** to save the settings.

A message appears that explains that if a null-modem cable that has no carrier detect pin wired is being used, the WAE may appear unresponsive on the console until the carrier detect signal is asserted. To recover from a misconfiguration, the WAE should be rebooted and the 0x2000 bootflag should be set to ignore the carrier detect setting.

Step 5 Click **OK** to continue.

To configure console line carrier detection from the CLI, you can use the **line console carrier-detect** global configuration command.

Configuring Remote Authentication Server Settings for WAAS Devices

If you have determined that your login authentication scheme is to include one or more external authentication servers, you must configure these server settings before you can configure the authentication scheme in the WAAS Central Manager GUI. The section contains the following topics:

- Configuring RADIUS Server Authentication Settings, page 7-12
- About TACACS+ Server Authentication Settings, page 7-14
- Configuring TACACS+ Server Settings, page 7-15
- Configuring Windows Domain Server Authentication Settings, page 7-17
- LDAP Server Signing, page 7-23

Configuring RADIUS Server Authentication Settings

RADIUS is a client/server authentication and authorization access protocol used by a network access server (NAS) to authenticate users attempting to connect to a network device. The NAS functions as a client, passing user information to one or more RADIUS servers. The NAS permits or denies network access to a user based on the response that it receives from one or more RADIUS servers. RADIUS uses the User Datagram Protocol (UDP) for transport between the RADIUS client and server.

RADIUS authentication clients reside on devices that are running WAAS software. When enabled, these clients send authentication requests to a central RADIUS server, which contains user authentication and network service access information.

You can configure a RADIUS key on the client and server. If you configure a key on the client, it must be the same as the one configured on the RADIUS servers. The RADIUS clients and servers use the key to encrypt all RADIUS packets transmitted. If you do not configure a RADIUS key, packets are not encrypted. The key itself is never transmitted over the network.

Note For more information about how the RADIUS protocol operates, see RFC 2138, *Remote Authentication Dial In User Service (RADIUS)*.

RADIUS authentication usually occurs when an administrator first logs in to the WAAS device to configure the device for monitoring, configuration, or troubleshooting purposes. RADIUS authentication is disabled by default. You can enable RADIUS authentication and other authentication methods at the same time. You can also specify which method to use first.

You can configure multiple RADIUS servers; authentication is attempted on the servers in order. If the first server is unreachable, then authentication is attempted on the other servers in the farm, in order. If authentication fails for any reason other than a server is unreachable, authentication is not attempted on the other servers in the farm.

Tip The WAAS Central Manager does not cache user authentication information. Therefore, the user is reauthenticated against the RADIUS server for every request. To prevent performance degradation caused by many authentication requests, install the WAAS Central Manager device in the same location as the RADIUS server, or as close as possible to it, to ensure that authentication requests can occur as quickly as possible.

To centrally configure RADIUS server settings for a WAAS device or device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Security > AAA > RADIUS**. The RADIUS Server Settings window appears. (See Figure 7-3.)

Figure 7-3 RADIUS Server Settings Window

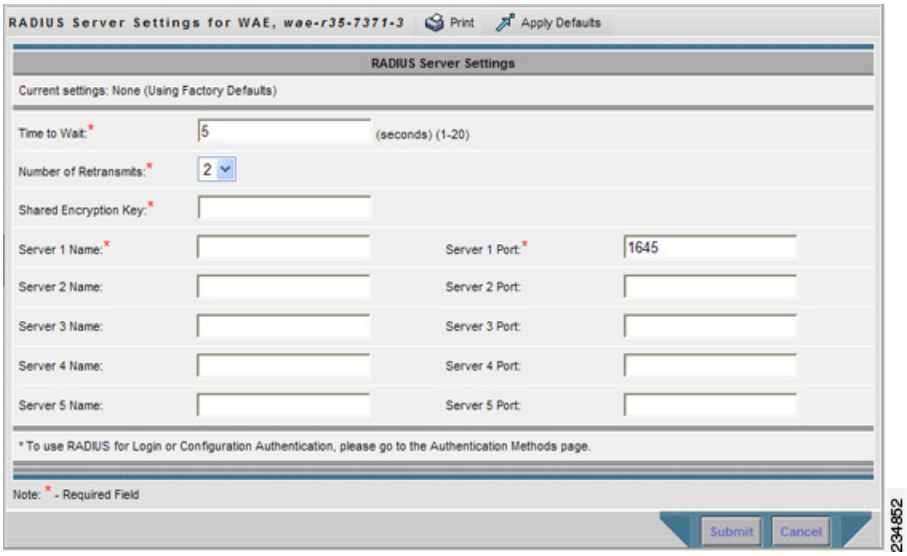

Step 3 In the Time to Wait field, specify how long the device or device group should wait for a response from the RADIUS server before timing out. The range is from 1 to 20 seconds. The default value is 5 seconds.

Step 4 In the Number of Retransmits field, specify the number of attempts allowed to connect to a RADIUS server. The default value is 2 times.

Step 5 In the Shared Encryption Key field, enter the secret key that is used to communicate with the RADIUS server.

Note If you configure a RADIUS key on the WAAS device (the RADIUS client), make sure that you configure an identical key on the external RADIUS server. Do not use the following characters: space, backwards single quote (`), double quote ("), pipe (|), or question mark (?).

Step 6 In the Server Name field, enter an IP address or hostname of the RADIUS server. Five different hosts are allowed.

Step 7 In the Server Port field, enter a UDP port number on which the RADIUS server is listening. You must specify at least one port. Five different ports are allowed.

Step 8 Click **Submit** to save the settings.

You can now enable RADIUS as an administrative login authentication and authorization method for this WAAS device or device group, as described in the "Enabling Administrative Login Authentication and Authorization Schemes for WAAS Devices" section on page 7-26.

To configure RADIUS settings from the CLI, you can use the **radius-server** global configuration command.

About TACACS+ Server Authentication Settings

TACACS+ controls access to network devices by exchanging network access server (NAS) information between a network device and a centralized database to determine the identity of a user or an entity. TACACS+ is an enhanced version of TACACS, a UDP-based access-control protocol specified by RFC 1492. TACACS+ uses TCP to ensure reliable delivery and encrypt all traffic between the TACACS+ server and the TACACS+ daemon on a network device.

TACACS+ works with many authentication types, including fixed password, one-time password, and challenge-response authentication. TACACS+ authentication usually occurs when an administrator first logs in to the WAAS device to configure the WAE for monitoring, configuring, or troubleshooting.

When a user requests restricted services, TACACS+ encrypts the user password information using the MD5 encryption algorithm and adds a TACACS+ packet header. This header information identifies the packet type being sent (for example, an authentication packet), the packet sequence number, the encryption type used, and the total packet length. The TACACS+ protocol then forwards the packet to the TACACS+ server.

A TACACS+ server can provide authentication, authorization, and accounting functions. These services, while all part of TACACS+, are independent of one another, so a given TACACS+ configuration can use any or all of the three services.

When the TACACS+ server receives a packet, it does the following:

- Authenticates the user information and notifies the client that the login authentication has either succeeded or failed.
- Notifies the client that authentication will continue and that the client must provide additional information. This challenge-response process can continue through multiple iterations until login authentication either succeeds or fails.

You can configure a TACACS+ key on the client and server. If you configure a key on a WAAS device, it must be the same as the one configured on the TACACS+ servers. The TACACS+ clients and servers use the key to encrypt all TACACS+ packets transmitted. If you do not configure a TACACS+ key, packets are not encrypted.

TACACS+ authentication is disabled by default. You can enable TACACS+ authentication and local authentication at the same time.

You can configure one primary and two backup TACACS+ servers; authentication is attempted on the primary server first. If the primary server is unreachable, then authentication is attempted on the other servers in the farm, in order. If authentication fails for any reason other than a server is unreachable, authentication is not attempted on the other servers in the farm.

The TACACS+ database validates users before they gain access to a WAAS device. TACACS+ is derived from the United States Department of Defense (RFC 1492) and is used by Cisco Systems as an additional control of nonprivileged and privileged mode access. The WAAS software supports TACACS+ only and not TACACS or Extended TACACS.

If you are using TACACS+ for user authentication, you can create WAAS user group names that match the user groups that you have defined on the TACACS+ server. WAAS can then dynamically assign roles and domains to users based on their membership in the groups defined on the TACACS+ server. (See the "Working with Accounts" section on page 8-3.) You must specify associated group names for each user in the TACACS+ configuration file, as follows:

```
user = tacusr1 {
 default service = permit
 service = exec
 {
   waas_rbac_groups = admin,groupname1,groupname2
   priv-lvl = 15
 }
 global = cleartext "tac"
}
```

For each user, list the groups they belong to in the waas_rbac_groups attribute, separating each group from the next with a comma.

The dynamic assignment of roles and domains based on external user groups requires a TACACS+ server that supports shell custom attributes. For example, these are supported in Cisco ACS 4.x and 5.1 and later.

Tip The WAAS Central Manager does not cache user authentication information, so the user is reauthenticated against the TACACS+ server for every request. To prevent performance degradation caused by many authentication requests, install the WAAS Central Manager device in the same location as the TACACS+ server, or as close as possible to it, to ensure that authentication requests can occur as quickly as possible.

Configuring TACACS+ Server Settings

The WAAS software CLI EXEC mode allows you to set, view, and test system operations. The mode is divided into two access levels: user and privileged. To access privileged-level EXEC mode, enter the **enable** EXEC command at the user access level prompt and specify the admin password when prompted for a password.

In TACACS+, the enable password feature allows an administrator to define a different enable password per administrative-level user. If an administrative-level user logs in to the WAAS device with a normal-level user account (privilege level of 0) instead of an admin or admin-equivalent user account (privilege level of 15), that user must enter the admin password to access privileged-level EXEC mode.

```
WAE> enable
Password:
```

Note This caveat applies even if the WAAS users are using TACACS+ for login authentication.

To centrally configure TACACS+ server settings on a WAAS device or device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Security > AAA > TACACS+**. The TACACS+ Server Settings window appears.

Note The TACACS+ server configuration cannot be modified or deleted when AAA Command Authorization is enabled.

Step 3 Check the **Use ASCII Password Authentication** check box to use the ASCII password type for authentication.

The default password type is PAP (Password Authentication Protocol). However, you can change the password type to ASCII when the authentication packets are to be sent in ASCII cleartext format.

Step 4 In the Time to Wait field, specify how long the device should wait before timing out. The range is from 1 to 20 seconds. The default value is 5 seconds.

Step 5 In the Number of Retransmits field, specify the number of attempts allowed to connect to a TACACS+ server. The range is 1 to 3 times. The default value is 2 times.

Step 6 In the Security Word field, enter the secret key that is used to communicate with the TACACS+ server.

Note If you configure a TACACS+ key on the WAAS device (the TACACS+ client), make sure that you configure an identical key on the external TACACS+ server. Do not use the following characters: space, backwards single quote (`), double quote ("), pipe (|), number sign (#), question mark (?), or backslash (\). The key is limited to 32 characters.

Step 7 In the Primary Server field, enter an IP address or hostname for the primary TACACS+ server.

If you want to change the default port (49), enter the port in the Primary Server Port field.

Step 8 In the Secondary Server field, enter an IP address or hostname for a secondary TACACS+ server.

If you want to change the default port (49), enter the port in the Secondary Server Port field.

Step 9 In the Tertiary Server field, enter an IP address or hostname for a tertiary TACACS+ server.

If you want to change the default port (49), enter the port in the Tertiary Server Port field.

Note You can specify up to two backup TACACS+ servers.

Step 10 Click **Submit** to save the settings.

You can now enable TACACS+ as an administrative login authentication and authorization method for this WAAS device or device group, as described in the "Enabling Administrative Login Authentication and Authorization Schemes for WAAS Devices" section on page 7-26.

To configure TACACS+ settings from the CLI, you can use the **tacacs** global configuration command.

Configuring Windows Domain Server Authentication Settings

A Windows domain controller can be configured to control access to the WAAS software services using either a challenge/response or shared secret authentication method. The system administrator can log in to the WAAS device by using an FTP, SSH, or Telnet session, the console, or the WAAS Central Manager GUI with a single user account (username/password/privilege). RADIUS and TACACS+ authentication schemes can be configured simultaneously with Windows domain authentication. Logging of a variety of authentication login statistics can be configured when Windows domain authentication is enabled. The log files and the statistical counters and related information can be cleared at any time.

In a WAAS network, Windows domain authentication is used in the following cases:

- Log in to the WAAS Central Manager GUI
- Log in to the WAE Device Manager GUI
- CLI configuration on any WAAS device

You can configure Windows authentication for the WAAS Central Manager device, a single WAAS device, or a group of devices. To configure Windows domain authentication on a WAAS device, you must configure a set of Windows domain authentication settings.

Note Windows domain authentication is not performed unless a Windows domain server is configured on the WAAS device. If the device is not successfully registered, authentication and authorization do not occur. WAAS supports authentication by a Windows domain controller running only on Windows Server 2000, Windows Server 2003, or Windows Server 2008.

If you are using NTLM authentication, the Windows domain server must be installed with the option to support pre-Windows 2000 operating systems. (On the installation Permissions screen of the Windows server dcpromo wizard, select “Permissions compatible with pre-Windows 2000 server operating systems.”)

This section contains the following topics:

Configuring Windows Domain Server Settings on a WAAS Device

You will need to know the name and IP address, or hostname, of the Windows domain controller that will be used for authentication.

Note If the Central Manager is version 4.2.3a or later and you want to configure the Windows domain settings on a WAAS device that is running version 4.2.3 or 4.2.1, you cannot use the Windows Domain Server Settings page on the Central Manager. You must use the **windows-domain diagnostics net** CLI command as described following the procedure below.

To configure Windows Domain server settings on a WAAS device or device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Security > AAA > Windows User Authentication**. The Windows User Authentication window appears. (See Figure 7-4.)

Note Workgroup settings are only required for Windows domain authentication, not for a domain join. You may skip to workgroup settings if you are only preforming a domain join.

Figure 7-4 Windows User Authentication

Step 3 In the Windows group for authorizing normal users field, specify an administrative group for normal users (non-superuser administrators), who only have restricted access to the chosen device (or device group) because their administrator user account has a privilege level of 0.

Note By default, there are not predefined user groups for Windows domain authorization configured on a WAE.

Step 4 In the Windows group for authorizing super users field, specify an administrative group for superusers (superuser administrators), who have unrestricted access to the chosen device (or device group) because their administrator user account has a privilege level of 15.

Note In addition to configuring Windows domain administrative group on a WAE, you must configure the Windows domain administrative group on your Microsoft Windows 2000 or 2003 server. You must create a Windows Domain administrative superuser group and a normal user group. Make sure that the group scope for the superuser group is set to global, assign user member to newly created administrative group, and add the user account (for example, the winsuper user) to the Windows domain superuser group. For more information about how to configure the Windows domain administrative group on your Windows server, see your Microsoft documentation.

When a user attempts to access this WAE through a Telnet session, FTP, or SSH session, the WAE is now configured to use the Active Directory user database to authenticate a request for administrative access.

Step 5 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 6 Choose **Configure > Security > Windows Domain > Domain Settings**. The Windows Domain tab appears. (See Figure 7-5.)

Figure 7-5 Windows Domain Tab

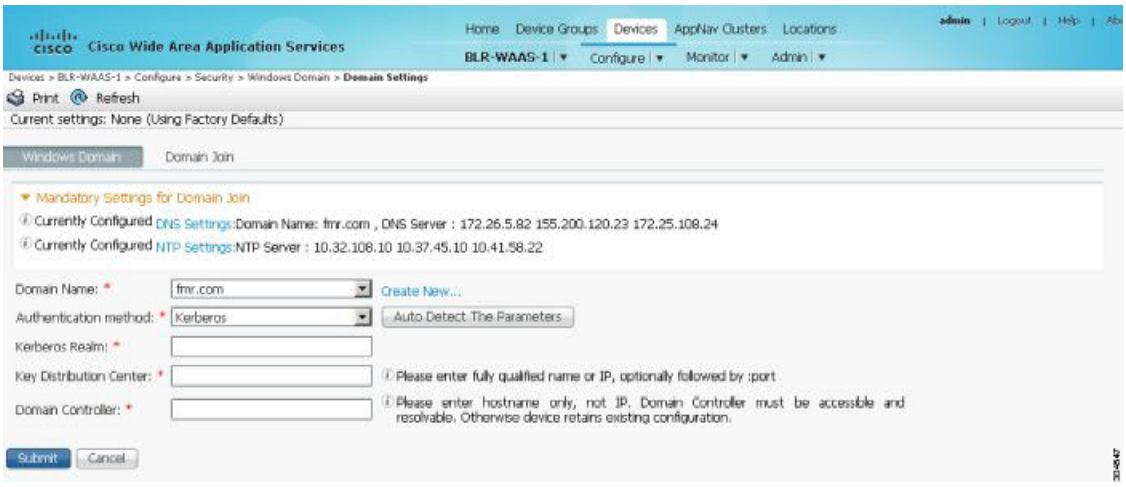

Note If the related WINS server and the workgroup or domain name have not been defined for the chosen device (or device group), an informational message is displayed at the top of this window to inform you that these related settings are currently not defined, as shown in Figure 7-5. To define these settings, choose **Configure > Network > WINS**.

Domain name, DNS server, and NTP configuration are mandatory prerequisites for the windows domain join. For full AAA functionality, workgroup and WINS server must also be configured.

NetBIOS name need not be configured for windows domain join. If left unconfigured, the first 15 characters of the hostname are automatically assigned as the NetBIOS name during the join.

Step 7 Select the Domain Name from the drop-down list or click **Create New** to create a new Local Domain Name.

Step 8 Select Kerberos or NTLM as a shared secure authentication method for administrative logins to the chosen device (or device group). The default authentication protocol is kerberos.

Note Windows domain user login authentication using NTLM protocol is deprecated in this release onwards. We recommend using Kerberos protocol for windows-domain user login authentication.

You must use Kerberos protocol for encrypted MAPI acceleration.

Click the **Auto Detect The Parameters** button when using kerberos to automatically obtain the kerberos realm, kerberos server, and domain controller. Domain, DNS, and NTP parameters must be configured first. This option is not supported with NTLM.

Once the device has been queried for the parameters, a status message will be displayed on the screen indicating either success or failure. The process may not be immediate and the status message will not appear until the auto detection process has been completed.

When successful, the parameters can be reviewed and edited, if required. Once the parameters have been reviewed, the values can then be submitted.

If the auto detection fails, you will need to check the configured domain/DNS configuration and enter them manually. The values can then be submitted.

Note Kerberos version 5 is used for Windows systems running Windows 2000 or higher with users logging in to domain accounts.

For Kerberos, skip the next step.

Step 9 For NTLM, select version 1 or version 2 from the drop-down list. NTLM version 1 is selected by default.

Note NTLM cannot be used for encrypted MAPI acceleration.

- NTLM version 1 is used for all Windows systems, including legacy systems such as Windows 98 with Active Directory, Windows NT, and more recent Windows systems, such as Windows 2000, Windows XP, and Windows 2003. We recommend the use of Kerberos if you are using a Windows 2000 SP4 or Windows 2003 domain controller.
- NTLM version 2 is used for Windows systems running Windows 98 with Active Directory, Windows NT 4.0 (Service Pack 4 or higher), Windows XP, Windows 2000, and Windows 2003. Enabling NTLM version 2 support on the WAAS print server will not allow access to clients who use NTLM or LM.

Caution Enable NTLM version 2 support in the print server only if all the clients' security policy has been set to Send NTLMv2 responses only/Refuse LM and NTLM.

Skip the next step.

Step 10 In the Kerberos Realm field, enter the fully qualified name of the realm in which the WAAS device resides. In the Key Distribution center, enter the fully qualified name or the IP address of the distribution center for the Kerberos key. If you clicked the **Auto Detect The Parameters** button when you selected Kerberos authentication method, these fields will already be populated.

All Windows 2000 domains are also Kerberos realms. Because the Windows 2000 domain name is also a DNS domain name, the Kerberos realm name for the Windows 2000 domain name is always in uppercase letters. This capitalization follows the recommendation for using DNS names as realm names in the Kerberos Version 5 protocol document (RFC-4120) and affects only interoperability with other Kerberos-based environments.

Step 11 In the Domain Controller field, enter the name of the Windows Domain Controller.

When you click **Submit**, the Central Manager validates this name by requesting the WAAS device (if version 4.2.x or later) to resolve the domain controller name. If the domain controller is not resolvable, you are asked to submit a valid name. If the device is offline, you are asked to verify device connectivity. If you are configuring a device group, the domain controller name is not validated on each device before this page is accepted and if it is not resolvable on a device, the configuration changes on this page are not applied to that device.

Step 12 Click **Submit**.

Note Make sure that you click **Submit** now so that the specified changes are committed to the WAAS Central Manager database. The Domain Administrator's username and password, which you will enter in Step 13, are not stored in the WAAS Central Manager's database.

Step 13 Register the chosen device (or device group) with the Windows Domain Controller as follows:

a. Click the **Domain Join** tab. (See Figure 7-6.)

Figure 7-6 Domain Join Tab

b. In the User Name field, enter a username (the domain\username or the domain name plus the username) for the specified Windows Domain Controller. This must be the username and password of a user who has administrative privileges in Active Directory (permission to add a computer to a domain).

For NTLM, the user credentials can be any normal user belonging to the Domain Users group. For Kerberos, the user credentials must be a user that belongs to the Domain Admins group, but need not be the system default Administrator user.

Note To use Windows domain server authentication, the WAAS device must join the Windows domain. For registration, you will need a user credential with permission to join a machine to the Windows domain. The user credential used for registration is not shown in clear text anywhere, including log files. WAAS does not modify the structure or schema of Windows Active Directory.

Note A domain join is required for encrypted MAPI acceleration using a machine account.

c. In the Password field, enter the password of the specified Windows Domain Controller account.

d. In the Confirm password field, reenter the password of the specified Windows Domain Controller.

e. (Optional) If desired, enter the name of the organizational unit in the Organizational Unit field (for Kerberos authentication only).

f. Click the **Join** button.

Note When you click the Join button, the WAAS Central Manager immediately sends a registration request to the WAAS device (or all of the devices in the device group) using SSH (the specified domain administrator password is encrypted by SSH). The registration request instructs the device to perform domain registration with the specified Windows Domain Controller using the specified domain username and password. If the device is accessible (if it is behind a NAT and has an external IP address), the registration request is performed by the device (or device group).

g. To check the status of the registration request, click the **Show Join Status** button.

   The status of domain join for the device (or all of the devices in the device group) is shown. It may take a few moments for the results to be updated.

h. If the join request fails, the result is shown in the join status window. Wait a few more minutes and try again to see the updated authentication status.

   If the request succeeds, the domain registration status is shown in the Domain Join Status window.

After configuring the Windows domain settings, to complete the process of enabling Windows authentication, you must set Windows as the authentication and authorization method for the device by using the Authentication Methods window, as described in the "Enabling Administrative Login Authentication and Authorization Schemes for WAAS Devices" section on page 7-26.

We recommend that you use the WAAS Central Manager GUI instead of the WAAS CLI to configure Windows Domain server settings, but if you want to use the CLI, see the following commands in the *Cisco Wide Area Application Services Command Reference*: **windows-domain join** and **kerberos** (if you are using Kerberos as a shared secure authentication method).

Next, register the WAAS device with the Windows domain server that you configured, by using the following command:

```
WAE# windows-domain join domain-name DomainName user UserName password Password
```

Finally, enable Windows Domain as the administrative login authentication and authorization configuration by using the following commands:

```
WAE(config)# authentication login windows-domain enable primary
WAE(config)# authentication configuration windows-domain enable primary
```

Unregistering a WAE from a Windows Domain Controller

If you want to unregister a WAE device from a Windows domain controller, you can do that directly from the WAAS Central Manager, as long as you have used the Kerberos shared secure authentication method. If you have used the NTLM method, you cannot unregister the WAE by using the WAAS Central Manager; you must log into the domain controller and remove the device registration manually.

Before you can unregister a device, you must disable windows authentication for the device.

To unregister a WAE device, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-name*).

Step 2 Choose **Configure > Security > AAA > Authentication Methods**. The Authentication and Authorization Methods window appears. (See Figure 7-7 on page 7-28.)

Step 3 Under both the Authentication Login Methods and the Authorization Methods sections, change each of the drop-down lists that are set to WINDOWS so that they are set to something different. For more information about changing these settings, see the "Enabling Administrative Login Authentication and Authorization Schemes for WAAS Devices" section on page 7-26.

Step 4 Click **Submit** to save the settings.

Step 5 Choose **Configure > Security > Windows Domain > Domain Settings** and click the **Domain Join** tab.

Step 6 (Optional) Enter the administrative username and password in the Administrator Username, Password, and Confirm Password fields. It is not mandatory to enter the username and password, but in some cases, the domain controller requires them to perform the unregistration.

Step 7 Click the **Leave** button.

Note When you click the Leave button, the WAAS Central Manager immediately sends an unregistration request to the WAAS device (or device group) using SSH. The unregistration request instructs the device to unregister from the specified Windows Domain Controller.

Request to unregister the device is not allowed when encrypted MAPI is configured to use machine accounts. You must delete the machine account identity before proceeding with the leave.

Step 8 Check the status of the unregistration request by waiting a few minutes and clicking the **Show Join Status** button.

If you want to use the CLI to unregister a WAE device, you must first use the following commands to disable windows authentication:

```
WAE(config)# no authentication login windows-domain enable
WAE(config)# no authentication configuration windows-domain enable
```

Next, unregister the WAAS device from the Windows domain server by using the following command (for Kerberos authentication):

```
WAE# windows-domain leave user UserName password Password
```

There is no CLI command to unregister the WAAS device if it is using NTLM authentication.

LDAP Server Signing

LDAP server signing is a configuration option of the Microsoft Windows Server's Network security settings. This option controls the signing requirements for Lightweight Directory Access Protocol (LDAP) clients. LDAP signing is used to verify that an intermediate party did not tamper with the LDAP packets on the network and to guarantee that the packaged data comes from a known source. Windows Server 2003 administration tools use LDAP signing to secure communications between running instances of these tools and the servers that they administer.

By using the Transport Layer Security (TLS, RFC 2830) protocol to provide communications privacy over the Internet, client/server applications can communicate in a way that prevents eavesdropping, tampering, or message forging. TLS v1 is similar to Secure Sockets Layer (SSL). TLS offers the same encryption on regular LDAP connections (ldap://:389) as SSL, while operating on a secure connection (ldaps://:636). A server certificate is used by the TLS protocol to provide a secure, encrypted connection to the LDAP server. A client certificate and key pair are required for client authentication.

In the WAAS software, login authentication with Windows 2003 domains is supported when the *LDAP server signing requirements* option for the Domain Security Policy is set to "Require signing." The LDAP server signing feature allows the WAE to join the domain and authenticate users securely.

Note When you configure your Windows domain controller to require an LDAP signature, you must also configure LDAP signing on the client WAE. By not configuring the client to use LDAP signatures, communication with the server is affected, and user authentication, group policy settings, and logon

scripts might fail. Install the Certification Authority service on the Microsoft server with the server's certificate (**Programs > Administrative Tools > Certification Authority**). Enable the LDAP server signing requirements property on the Microsoft server (**Start > Programs > Administrative Tools > Domain Controller Security Policy**). In the displayed window, choose **Require signing** from the drop-down list, and click **OK**.

For information about how to configure your Windows domain controller to require an LDAP signature, see your Microsoft documentation.

This section contains the following topics:

- Configuring LDAP Signing on the Client WAEs, page 7-24
- Disabling LDAP Server Signing on a Client WAE, page 7-25

Configuring LDAP Signing on the Client WAEs

You can configure a security setting on Windows 2003 domain controllers to require clients (such as WAEs) to sign LDAP requests. Because unsigned network traffic can be intercepted and manipulated by outside parties, some organizations require LDAP server signing to prevent man-in-the-middle attacks on their LDAP servers. You can only configure LDAP signing on a single WAE; it cannot be configured at a system level. In addition, you must configure LDAP signing on a WAE through the WAAS CLI; you cannot configure LDAP signing through any of the WAAS GUIs (either the WAAS Central Manager GUI or the WAE Device Manager GUI).

By default, LDAP server signing is disabled on a WAE. To enable this feature on a WAE, follow these steps:

Step 1 Enable LDAP server signing on the WAE:

```
WAE# configure
WAE(config)# smb-conf section "global" name "ldap ssl" value "yes"
```

Step 2 Save the configuration on the WAE:

```
WAE(config)# exit
WAE# copy run start
```

Step 3 Check the current running LDAP client configuration on the WAE:

```
WAE# show smb-conf
```

Step 4 Register the WAE with the Windows domain:

```
WAE# windows-domain diagnostics net "ads join -U username%password"
```

Step 5 Enable user login authentication on the WAE:

```
WAE# configure
WAE(config)# authentication login windows-domain enable primary
```

Step 6 Enable user login authorization on the WAE:

```
WAE(config)# authentication configuration windows-domain enable primary
```

Step 7 Check the current configuration for login authentication and authorization on the WAE:

```
WAE# show authentication user
Login Authentication:    Console/Telnet/Ftp/SSH Session
---------------------------- ------------------------------
local                                 enabled (secondary)
```

```
Windows domain           enabled (primary)
Radius                            disabled
Tacacs+                          disabled

Configuration Authentication: Console/Telnet/Ftp/SSH Session
---------------------------- -----------------------------
local                                 enabled (primary)
Windows domain           enabled (primary)
Radius                             disabled
Tacacs+                           disabled
```

The WAE is now configured to authenticate Active Directory users. Active Directory users can use Telnet, FTP, or SSH to connect to the WAE or they can access the WAE through the WAAS GUIs (WAAS Central Manager GUI or the WAE Device Manager GUI).

Step 8 View statistics that are related to Windows domain user authentication. Statistics increment after each user authentication attempt:

```
WAE# show statistics windows-domain
Windows Domain Statistics
  ----------------------------------------------
  Authentication:
    Number of access requests:                     9
    Number of access deny responses:               3
    Number of access allow responses:              6
  Authorization:
    Number of authorization requests:               9
    Number of authorization failure responses:      3
    Number of authorization success responses:      6
  Accounting:
    Number of accounting requests:                  0
    Number of accounting failure responses:         0
    Number of accounting success responses:         0

WAE# show statistics authentication
    Authentication  Statistics
    -------------------------------------
    Number of access requests:        9
    Number of access deny responses:  3
    Number of access allow responses: 6
```

Step 9 Use the **clear statistics** EXEC command to clear the statistics on the WAE:

- To clear all of the login authentication statistics, enter the **clear statistics authentication** EXEC command.
- To clear only the statistics that are related to Windows domain authentication, enter the **clear statistics windows-domain** EXEC command.
- To clear all of the statistics, enter the **clear statistics all** EXEC command.

Disabling LDAP Server Signing on a Client WAE

To disable LDAP server signing on a WAE, follow these steps:

Step 1 Unregister the WAE from the Windows domain:

```
WAE# windows-domain diagnostics net "ads leave -U Administrator"
```

Step 2 Disable user login authentication:

```
WAE# configure
WAE(config)# no authentication login windows-domain enable primary
```

Step 3 Disable LDAP signing on the WAE:

```
WAE(config)# no smb-conf section "global" name "ldap ssl" value "yes"
```

Enabling Administrative Login Authentication and Authorization Schemes for WAAS Devices

This section describes how to centrally enable the various administrative login authentication and authorization schemes (the authentication configuration) for a WAAS device or device group.

Caution Make sure that RADIUS, TACACS+, or Windows domain authentication is configured and operating correctly before disabling local authentication and authorization. If you disable local authentication and if RADIUS, TACACS+, or Windows domain authentication is not configured correctly, or if the RADIUS, TACACS+, or Windows domain server is not online, you may be unable to log in to the WAAS device.

By default, a WAAS device uses the local database to authenticate and authorize administrative login requests. The WAAS device verifies whether all authentication databases are disabled and if so, sets the system to the default state. For information on this default state, see the "Default Administrative Login Authentication and Authorization Configuration" section on page 7-4.

Note You must configure the TACACS+, or RADIUS, or Windows server settings for the WAAS device (or device group) before you configure and submit these settings. See the "About TACACS+ Server Authentication Settings" section on page 7-14, the "Configuring RADIUS Server Authentication Settings" section on page 7-12, and the "Configuring Windows Domain Server Authentication Settings" section on page 7-17 for information on how to configure these server settings on a WAAS device or device group.

By default, WAAS devices fail over to the secondary method of administrative login authentication whenever the primary administrative login authentication method fails for any reason. You change this default login authentication failover method through the WAAS Central Manager GUI, as follows:

- To change the default for a WAAS device, choose **Devices >** *device-name* and then choose **Configure > Security > AAA > Authentication Methods** from the menu. Check the **Failover to next available authentication method** box in the displayed window and click **Submit**.
- To change the default for a device group, choose **Device Groups >** *device-group-name* and then choose **Configure > Security > AAA > Authentication Methods** from the menu. Check the **Failover to next available authentication method** box in the displayed window and click **Submit**.

After you enable the failover to next available authentication method option, the WAAS device (or the devices in the device group) queries the next authentication method only if the administrative login authentication server is unreachable, not if authentication fails for some other reason. The authentication server could be unreachable due to an incorrect key in the RADIUS or TACACS+ settings on the WAAS device.

You can configure multiple TACACS+ or RADIUS servers; authentication is attempted on the primary server first. If the primary server is unreachable, then authentication is attempted on the other servers in the TACACS+ or RADIUS farm, in order. If authentication fails for any reason other than a server is unreachable, authentication is not attempted on the other servers in the farm. This process applies regardless of the setting of the **Failover to next available authentication method** check box.

Note To use the login authentication failover feature, you must set TACACS+, RADIUS, or Windows domain as the primary login authentication method, and local as the secondary login authentication method.

If the failover to next available authentication method option is *enabled*, follow these guidelines:

- You can configure only two login authentication schemes (a primary and secondary scheme) on the WAAS device.
- Note that the WAAS device (or the devices in the device group) fails over from the primary authentication scheme to the secondary authentication scheme only if the specified authentication server is unreachable.
- Configure the local database scheme as the secondary scheme for both authentication and authorization (configuration).

For example, if the failover to next available authentication method option is enabled and RADIUS is set as the primary login authentication scheme and local is set as the secondary login authentication scheme, the following events occur:

1. When the WAAS device (or the devices in the device group) receives an administrative login request, it queries the external RADIUS authentication server.
2. One of the following occurs:
   a. If the RADIUS server is reachable, the WAAS device (or the devices in the device group) uses this RADIUS database to authenticate the administrator.
   b. If the RADIUS server is not reachable, the WAAS device (or the devices in the device group) tries the secondary authentication scheme (that is, it queries its local authentication database) to authenticate the administrator.

Note The local database is contacted for authentication only if this RADIUS server is not available. In any other situation (for example, if the authentication fails in the RADIUS server), the local database is not contacted for authentication.

Conversely, if the failover to next available authentication method option is *disabled*, then the WAAS device (or the devices in the device group) contacts the secondary authentication database regardless of the reason why the authentication failed with the primary authentication database.

If all the authentication databases are enabled for use, then all the databases are queried in the order of priority selected and based on the failover reason. If no failover reason is specified, then all the databases are queried in the order of their priority. For example, first the primary authentication database is queried, then the secondary authentication database is queried, then the tertiary database is queried, and finally the quaternary authentication database is queried.

To specify the login authentication and authorization scheme for a WAAS device or device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Security > AAA > Authentication Methods**. The Authentication and Authorization Methods window appears. (See Figure 7-7.)

Figure 7-7 Authentication and Authorization Methods Window

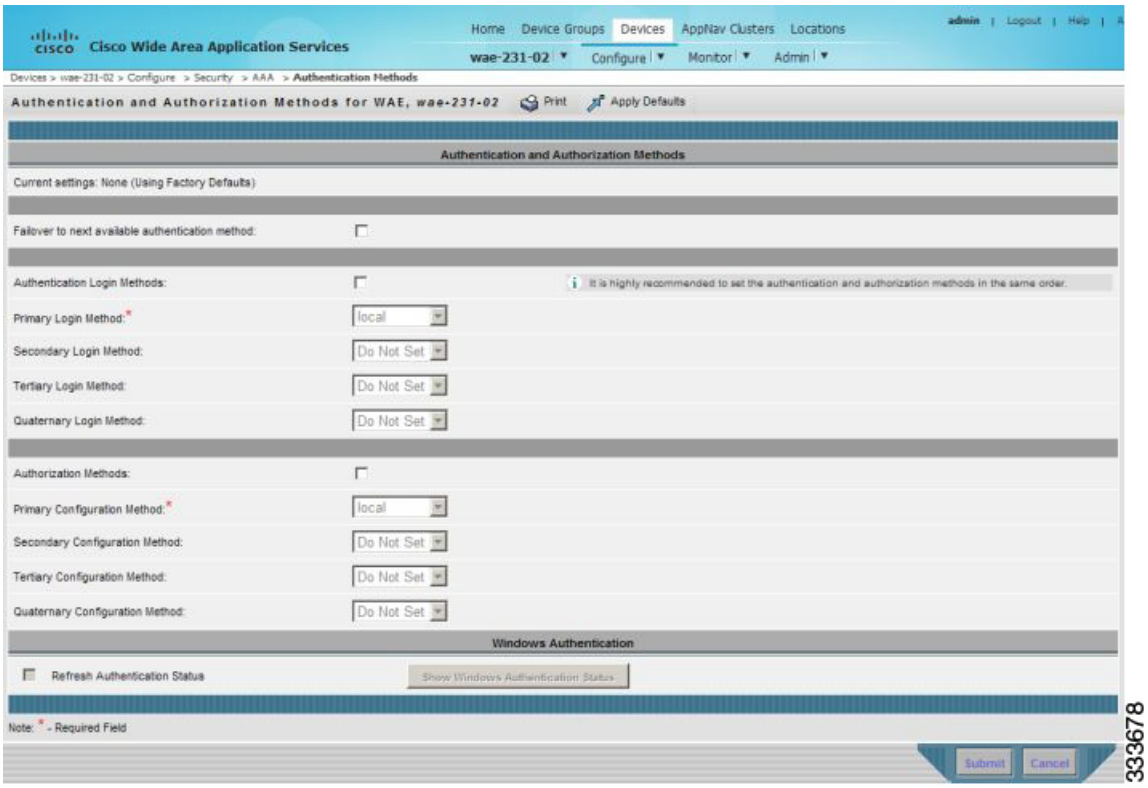

Step 3 Check the **Failover to next available authentication method** check box to query the secondary authentication database only if the primary authentication server is unreachable. When the box is unchecked, the other authentication methods are tried if the primary method fails for any reason.

To use this feature, you must set TACACS+, RADIUS, or Windows domain as the primary authentication method and local as a secondary authentication method. Make sure that you configure the local method as a secondary scheme for both authentication and authorization (configuration).

Step 4 Check the **Authentication Login Methods** check box to enable authentication privileges using the local, TACACS+, RADIUS, or WINDOWS databases.

Step 5 Specify the order of the login authentication methods that the chosen device or device group are to use:

a. From the Primary Login Method drop-down list, choose **local**, **TACACS+**, **RADIUS**, or **WINDOWS**. This option specifies the first method that the chosen device (or the device group) should use for administrative login authentication.

b. From the Secondary Login Method drop-down list, choose **local**, **TACACS+**, **RADIUS**, or **WINDOWS**. This option specifies the method that the chosen device (or the device group) should use for administrative login authentication if the primary method fails.

c. From the Tertiary Login Method drop-down list, choose **local**, **TACACS+**, **RADIUS**, or **WINDOWS**. This option specifies the method that the chosen device (or the device group) should use for administrative login authentication if both the primary and the secondary methods fail.

d. From the Quaternary Login Method drop-down list, choose **local**, **TACACS+**, **RADIUS**, or **WINDOWS**. This option specifies the method that the chosen device (or device group) should use for administrative login authentication if the primary, secondary, and tertiary methods all fail.

Note We strongly recommend that you specify the local method as the last method in your prioritized list of login authentication and authorization methods. By adhering to this practice, the WAAS administrator will be able to still log in to a WAAS device (or the devices in the device groups) through the local authentication and authorization method if the specified external third-party servers (TACACS+, RADIUS, or Windows domain servers) are not reachable.

Step 6 Check the **Authorization Methods** check box to enable authorization privileges using the local, TACACS+, RADIUS, or WINDOWS databases.

Note Authorization privileges apply to console and Telnet connection attempts, secure FTP (SFTP) sessions, and Secure Shell (SSH, Version 1 and Version 2) sessions.

Step 7 Specify the order of the login authorization (configuration) methods that the chosen device (or the device group) should use:

Note We strongly recommend that you set the administrative login authentication and authorization methods in the same order. For example, configure the WAAS device (or device group) to use RADIUS as the primary login method, TACACS+ as the secondary login method, Windows as the tertiary method, and the local method as the quaternary method for both administrative login authentication and authorization.

a. From the Primary Configuration Method drop-down list, choose **local**, **TACACS+**, **RADIUS**, or **WINDOWS**. This option specifies the first method that the chosen device (or the device group) should use to determine authorization privileges.

Note If you have checked the **Failover to next available authentication method** check box (Step 3), make sure that you choose **TACACS+ or RADIUS** from the Primary Configuration Method drop-down list to configure either the TACACS+ or RADIUS method as the primary scheme for authorization (configuration).

b. From the Secondary Configuration Method drop-down list, choose **local**, **TACACS+**, **RADIUS**, or **WINDOWS**. This option specifies the method that the chosen device (or the device group) should use to determine authorization privileges if the primary method fails.

Note If you have checked the **Failover to next available authentication method** check box (Step 3), make sure that you choose **local** from the Secondary Configuration Method drop-down list to configure the local method as the secondary scheme for authorization (configuration).

c. From the Tertiary Configuration Method drop-down list, choose **local**, **TACACS+**, **RADIUS**, or **WINDOWS**. This option specifies the method that the chosen device (or the device group) should use to determine authorization privileges if both the primary and secondary methods fail.

d. From the Quaternary Configuration Method drop-down list, choose **local**, **TACACS+**, **RADIUS**, or **WINDOWS**. This option specifies the method that the chosen device (or device group) should use to determine authorization privileges if the primary, secondary, and tertiary methods all fail.

Step 8 To refresh the authentication status, check the box and click the **Show Windows Authentication Status** button. This option is only available when Windows is set as the authentication and authorization methods.

A dialog box appears prompting if you want to continue with this request to refresh the status of the authentication request. (See Figure 7-8.)

Figure 7-8 Confirmation Dialog Box

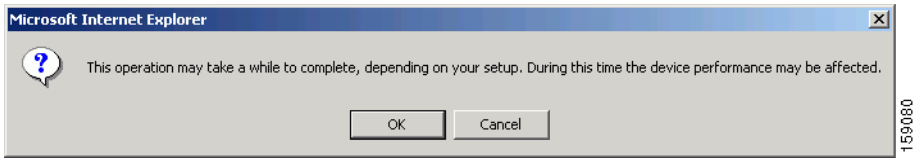

Click **OK** to continue or **Cancel** to cancel the request.

If the request fails, you will receive an error dialog. Wait a few more minutes and try again to see the updated authentication status.

Step 9 Click **Submit** to save the settings.

Note If you have enabled the Windows authentication or authorization method, the Central Manager queries the WAE (of version 4.2.1 or higher) to ensure that it is registered to a Windows domain. This can take up to one minute after you click **Submit**. You will see a message asking you to confirm this process and you must click **OK** to proceed. If you are configuring a WAE of version 4.1.x or lower, or a device group, the Central Manager does not query the WAE(s) and you must ensure that each WAE is properly registered. You will see a message informing you that system behavior is unknown (if a WAE is unregistered) and you must click **OK** to proceed.

Note If you have enabled the Windows authentication method, it takes about 15 seconds to activate it. Wait at least 15 seconds before checking Windows authentication status or performing any operation that requires Windows authentication.

To configure the login authentication and authorization scheme from the CLI, you can use the **authentication** global configuration command. Before you can enable Windows domain authentication or authorization for a device, the device must be registered with the Windows domain controller.

Configuring AAA Command Authorization

Command authorization enforces authorization through an external AAA server for each command executed by the CLI user. All commands executed by a CLI user are authorized before they are executed. RADIUS, Windows domain, and local users are not affected.

Note Only commands executed through the CLI interface are subject to command authorization.

When command authorization is enabled, you must specify "permit null" on the TACACS+ server to allow authorized commands with no arguments to be executed.

To configure command authorization for a WAAS device or device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Security > AAA > Command Authorization Settings**. The Command Authorization window appears.

Step 3 Check the Command Authorization Level check box for the desired level.

- Level 0—Only EXEC commands are authorized by the TACACS+ server before they are executed, regardless of user level (normal or super). Global configuration commands are not allowed.
- Level 15—Both EXEC and global configuration level commands are authorized by the TACACS+ server before they are executed, regardless of user level (normal or super).

Note You must have a TACACS+ server configured before you can configure command authorization.

Step 4 Click **Submit** to save the settings.

Configuring AAA Accounting for WAAS Devices

Accounting tracks all user actions and when the actions occurred. It can be used for an audit trail or for billing for connection time or resources used (bytes transferred). Accounting is disabled by default.

The WAAS accounting feature uses TACACS+ server logging. Accounting information is sent to the TACACS+ server only, not to the console or any other device. The syslog file on the WAAS device logs accounting events locally. The format of events stored in the syslog is different from the format of accounting messages.

The TACACS+ protocol allows effective communication of AAA information between WAAS devices and a central server. It uses TCP for reliable connections between clients and servers. WAAS devices send authentication and authorization requests, as well as accounting information to the TACACS+ server.

Note Before you can configure the AAA accounting settings for a WAAS device, you must first configure the TACACS+ server settings for the WAAS device. (See the "About TACACS+ Server Authentication Settings" section on page 7-14.)

Note If you enable AAA accounting for a device, we strongly recommended that you create an IP ACL condition in the first entry position permitting access to the TACACS+ servers to avoid delay while processing the commands. For information on IP ACLs, see Chapter 9, "Creating and Managing IP Access Control Lists for WAAS Devices."

To centrally configure AAA accounting settings for a WAAS device or device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Security > AAA > AAA Accounting**. The AAA Accounting Settings window appears.

Step 3 From the **System Events** drop-down list, choose a keyword to specify when the chosen device (or the device group) should track system-level events that are not associated with users, such as reloads, and to activate accounting for system events.

Step 4 From the **Exec Shell and Login/Logout Events** drop-down list, choose a keyword to specify when the chosen device (or the device group) should track EXEC shell and user login and logout events and to activate accounting for EXEC mode processes. Reports include username, date, start and stop times, and the WAAS device IP address.

Step 5 From the **Normal User Commands** drop-down list, choose a keyword to specify when the chosen device (or the device group) should track all the commands at the normal user privilege level (privilege level 0) and to activate accounting for all commands at the non-superuser administrative (normal user) level.

Step 6 From the **Administrative User Commands** drop-down list, choose a keyword to specify when the chosen device (or the device group) should track all commands at the superuser privilege level (privilege level 15) and to activate accounting for all commands at the superuser administrative user level.

Caution Before using the **wait-start** option, ensure that the WAAS device is configured with the TACACS+ server and is able to successfully contact the server. If the WAAS device cannot contact a configured TACACS+ server, it might become unresponsive.

Table 7-2 describes the event type options.

Table 7-2 Event Types for AAA Accounting

GUI Parameter	Function
Event Type Options	
stop-only	The WAAS device sends a stop record accounting notice at the end of the specified activity or event to the TACACS+ accounting server.
start-stop	The WAAS device sends a start record accounting notice at the beginning of an event and a stop record at the end of the event to the TACACS+ accounting server. The start accounting record is sent in the background. The requested user service begins regardless of whether or not the start accounting record was acknowledged by the TACACS+ accounting server.

Table 7-2 Event Types for AAA Accounting (continued)

GUI Parameter	Function
wait-start	The WAAS device sends both a start and a stop accounting record to the TACACS+ accounting server. However, the requested user service does not begin until the start accounting record is acknowledged. A stop accounting record is also sent.
Do Not Set	Accounting is disabled for the specified event.

Step 7 Check the **Enable CMS CLI Accounting** check box to enable AAA accounting to TACACS+ server.

Step 8 Click **Submit** to save the settings.

To configure AAA accounting settings from the CLI, you can use the **aaa accounting** global configuration command.

Viewing Audit Trail Logs

The WAAS Central Manager device logs user activity in the system. The only activities that are logged are those activities that change the WAAS network. For more information on viewing a record of user activity on your WAAS system, see the "Viewing the Audit Trail Log" section on page 17-57.

CHAPTER 8

Creating and Managing Administrator User Accounts and Groups

This chapter describes how to create user accounts and groups from the WAAS Central Manager GUI.

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

This chapter contains the following sections:

Overview of Administrator User Accounts

Your WAAS system comes with an administrator account already created that you can use to access the WAAS Central Manager GUI as well as the WAAS CLI. This account has a username of *admin* and a password of *default*. You can use the WAAS Central Manager GUI to change the password of this account.

If you want to create additional administrator user accounts, see Table 8-1 for a description of the two types of accounts you can create from the WAAS Central Manager GUI.

Table 8-1 Account Type Descriptions

Account Type	Description
Roles-based account	Allows you to create accounts that manage and configure specific WAAS services. For example, you may want to delegate the configuration of application acceleration to a specific administrator. In this case, you could create a roles-based account that only has access to the Acceleration pages in the WAAS Central Manager GUI. You can also create a roles-based account that only has access to the WAE Device Manager instead of the WAAS Central Manager GUI. And you can create a role-based account that also is a local user account. You create roles-based accounts from the Admin menu in the WAAS Central Manager GUI.
Local account	Provides CLI access to WAE devices and optionally allows users to access the WAE Device Manager GUI. A user with this account type can log into the WAAS Central Manager but they have the access rights assigned to the default account, which initially has access to no GUI functionality. We recommend that you create a local account if there is an administrator that only needs CLI access to WAE devices or to the WAE Device Manager GUI. You create local accounts in the same way as roles-based accounts, but you check the Local User check box when creating the account.

Creating and Managing User Accounts

This section contains the following topics:

Overview for Creating an Account

Table 8-2 provides an overview of the steps you must complete to create a new roles-based administrator account.

Table 8-2 Checklist for Creating a Roles-based Administrator Account

Task	Additional Information and Instructions
1. Create a new account.	Creates an account on the system with a specific username, password, and privilege level. For more information, see the "Creating a New Account" section on page 8-4.
2. Create a role for the new account.	Creates a role that specifies the services that an account can configure in your WAAS network. For more information, see the "Creating a New Role" section on page 8-10. If you are using an external authentication server, you can define matching user groups that automatically assign roles to users.

Table 8-2 Checklist for Creating a Roles-based Administrator Account (continued)

Task	Additional Information and Instructions
3. Assign the role to the new account.	Assigns the new role to the new account. For more information, see the "Assigning a Role to a User Account" section on page 8-12. If you are using an external authentication server, you can define matching user groups that automatically assign roles to users.
4. Create a domain.	Creates a domain that will specify the WAEs, device groups, or AppNav Clusters that the new account can manage. For more information, see the "Creating a New Domain" section on page 8-14.
5. Add an entity to the domain.	Adds one or more WAEs, device groups, or AppNav Clusters to the domain. For more information, see the "Adding an Entity to a Domain" section on page 8-15.
6. Assign a domain to a user account.	Assigns the domain to the new user account. For more information, see the "Assigning a Domain to a User Account" section on page 8-15. If you are using an external authentication server, you can define matching user groups that automatically assign domains to users.

Working with Accounts

When you create a user account, you enter information about the user such as the username, the name of the individual who owns the account, contact information, job title, and department. All user account information is stored in an internal database on the WAAS Central Manager.

Each user account can then be assigned to a role. A *role* defines which WAAS Central Manager GUI configuration pages the user can access and which services the user has authority to configure or modify. The WAAS Central Manager provides one predefined role, known as the admin role. The admin role has access to all services. A *domain* defines which entities in the network that the user can access and configure or modify. You can assign a user account to zero or more roles and to zero or more domains.

In addition to user accounts, you can create user groups if you are using external authentication of users on a TACACS+ or Windows domain server (not a RADIUS server). By creating user group names that match the user groups that you have defined on the external authentication server, WAAS can dynamically assign roles and domains to users based on their membership in a group as defined on the external authentication server. You do not need to define a role or domain for each user individually.

Two default user accounts are preconfigured in the WAAS Central Manager. The first account, called *admin*, is assigned the administrator role that allows access to all services and access to all entities in the system. This account cannot be deleted from the system, but it can be modified. Only the username and the role for this account are unchangeable. Only an account that has been assigned the admin role can create other admin-level accounts.

The second preconfigured user account is called *default*. Any user account that is authenticated but has not been registered in the WAAS Central Manager obtains the access rights (role) assigned to the default account. This account is configurable by an administrator, but it cannot be deleted nor its username changed. Initially, the default account has no access to GUI functionality because it has no roles defined, though it can log into the WAAS Central Manager GUI.

This section contains the following topics:

- Viewing User Accounts, page 8-8
- Unlocking User Accounts, page 8-8

Creating a New Account

The first step in setting up an account is to create the account by specifying a username and selecting whether a local CLI account is created at the same time. After the account is created, you can assign roles to the account that determine the WAAS services and devices that the account can manage and configure.

Table 8-3 describes the results of creating a local CLI user when setting up an account.

Table 8-3 Results of Creating a Local User

Action	Result
Creating a Local User	• The account can be used to access the WAAS CLI, WAAS Central Manager GUI (with the default role), and WAE Device Manager (if that option is selected). • Users can change their own passwords, and the password change will propagate to standby WAAS Central Managers. • The account is stored in the WAAS Central Manager database and is also propagated to the standby WAAS Central Managers.
Not Creating a Local User	• The user account is created in the primary and standby WAAS Central Manager management databases. • No user account is created in the CLI. Users will have to use another account to access the CLI. • The new account can be used to log in to the WAAS Central Manager GUI if an external authentication server is set. The user is assigned the roles defined for the default user (initially none). • Local users can change their passwords using the WAAS Central Manager GUI only if they have roles that allow access to the Admin > AAA section.

Note If a user account has been created from the CLI only, when you log in to the WAAS Central Manager GUI for the first time, the Centralized Management System (CMS) automatically creates a user account (with the same username as configured in the CLI) with default authorization and access control. An account created from the CLI initially will be unable to access any configuration pages in the WAAS Central Manager GUI. You must use an admin account to give the account created from the CLI the roles that it needs to perform configuration tasks from the WAAS Central Manager GUI.

To create a new account, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > Users**.

The User Accounts window displays all the user accounts on the system.

Step 2 Click the **Create New User Accounts** icon.

The Creating New User Account window appears.

Note This window can be accessed only by users with administrator-level privileges.

Step 3 In the Username field, enter the user account name.

Usernames are case sensitive and cannot contain characters other than letters, numbers, period, hyphen, and underscore.

Step 4 Complete the following steps to allow the user to access the WAE Device Manager GUI:

a. Check the **WAE Device Manager User** check box.

b. From the Device Manager Access drop-down list, choose one of the following options for Device Manager GUI access for this account:

- **Read Only**—Limits this user to read only access to the Device Manager GUI.
- **Read Write**—Allows this user to have read and write access to the Device Manager GUI.

Step 5 Complete the following steps to create a local CLI user account:

a. Check the **Local User** check box. See Table 8-3 on page 8-4 for information about the benefits of creating a local CLI user. A local user is created on all WAE devices.

Note Do not create a local user that has a username identical to a username defined in an external authentication server that is authorizing access to the WAAS device.

b. In the Password field, enter a password for the local user account, and reenter the same password in the Confirm Password field. Passwords are case-sensitive, must be 1 to 31 characters in length, and cannot contain the characters ' " | (apostrophe, double quote, or pipe) or any control characters.

c. From the CLI Privilege Level drop-down list, select one of the following options for the local user account:

- **0 (normal user)**—Limits the CLI commands this user can use to only user-level EXEC commands. This is the default value.
- **15 (super user)**—Allows this user to use privileged EXEC-level CLI commands, similar to the functions that a Central Manager GUI user with the admin role can perform.

Note The WAAS CLI EXEC mode is used for setting, viewing, and testing system operations. It is divided into two access levels: user and privileged. A local user who has "normal" privileges can only access the user-level EXEC CLI mode. A local user who has "superuser" privileges can access the privileged EXEC mode as well as all other modes (for example, configuration mode and interface mode) to perform any administrative task. For more information about the user-level and privileged EXEC modes and CLI commands, see the *Cisco Wide Area Application Services Command Reference*.

Step 6 (Optional) In the User Information fields, enter the following information about the user in the appropriate fields: first name, last name, phone number, e-mail address, job title, and department.

Step 7 (Optional) In the Comments field. enter any additional information about this account.

Step 8 Click **Submit**.

A Changes Submitted message appears at the bottom of the window.

Step 9 Assign roles to this new account as described in the "Working with Roles" section on page 8-9 and assign domains as described in the "Working with Domains" section on page 8-14.

Modifying and Deleting User Accounts

Note Modifying a user account from the CLI does not update the Centralized Management System (CMS) database and the change will not be reflected in the Central Manager GUI.

To modify an existing user account, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > Users**.

The User Accounts window appears.

Step 2 Click the **Edit** icon next to the user account that you want to modify.

The Modifying User Account window appears. You can delete or edit user accounts as follows:

Note This window can only be accessed by users with administrator-level privileges.

- To delete the user account, click the **Delete** icon in the taskbar, and then click **OK** to confirm the deletion.

  If the local user account was created using the WAAS Central Manager GUI, the corresponding user account is removed from the CLI and is also deleted from all standby WAAS Central Managers.

  Note Deleting a user account from the CLI does *not* disable the corresponding user account in the CMS database. Consequently, the user account remains active in the CMS database. User accounts created in the WAAS Central Manager GUI should always be deleted from the WAAS Central Manager GUI.

- To edit the user account, make the necessary changes to the username and account information, and click **Submit**.

Changing the Password for Your Own Account

If you are logged in to the WAAS Central Manager GUI, you can change your own account password if you meet the following requirements:

- Your account and password were created in the WAAS Central Manager GUI and not in the CLI.
- You are authorized to access the password window.

Note We do not recommend changing the local CLI user password from the CLI. Any changes to local CLI user passwords from the CLI are not updated in the management database and are not propagated to the standby WAAS Central Manager. Therefore, passwords in the management database will not match a new password configured in the CLI.

Note The advantage of initially setting passwords from the WAAS Central Manager GUI is that both the primary and the standby WAAS Central Managers will be synchronized, and GUI users will not have to access the CLI to change their password.

To change the password for your own account, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > Security > Password**.

The Changing Password for User Account window appears.

Step 2 In the New Password field, enter the changed password. Passwords are case sensitive, must be 1 to 31 characters in length, and cannot contain the characters ' " | (apostrophe, double quote, or pipe) or any control characters.

Step 3 In the Confirm New Password field, reenter the password for confirmation.

Step 4 Click **Submit**.

The message "Changes Submitted" appears at the bottom of the window confirming that your password has been changed.

When you change the password of an account by using the WAAS Central Manager GUI, it changes the password for all WAE devices managed by the Central Manager.

Changing the Password for Another Account

If you log into the WAAS Central Manager GUI using an account with admin privileges, you can change the password of any other account.

Note If you change a user password from the CLI, the password change applies only to the local device, will not be reflected in the Central Manager GUI, and is not propagated to any other devices.

To change the password for another account, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > Users**.

A list of roles-based user accounts appears.

Step 2 Click the **Edit** icon next to the account that needs a new password. The Modifying User Account window appears.

Step 3 In the Password field, enter the changed password. Passwords are case-sensitive, must be 1 to 31 characters in length, and cannot contain the characters ' " | (apostrophe, double quote, or pipe) or any control characters.

Step 4 In the Confirm Password field, reenter the password for confirmation.

Step 5 Click **Submit**.

The message "Changes Submitted" appears at the bottom of the window confirming that your password has been changed.

Viewing User Accounts

To view all user accounts, choose **Admin > AAA> Users** from the WAAS Central Manager GUI. The User Accounts window displays all the user accounts in the management database. From this window, you can also create new accounts as described in the "Creating a New Account" section on page 8-4.

To view user accounts for a specific device, choose **Devices >** *device-name* and then choose *device-name* > **Device Users** or **CM Users**, depending on the device mode. The Users for device window displays all the user accounts defined for the device.

If a user account is locked out on the device, you can unlock it from this window. Check the box next to the account and click the **Unlock** button.

To view the details for an account, click the **View** icon next to the account.

Unlocking User Accounts

When a user account is locked out, the user cannot log in to the WAAS device until an administrator unlocks the account. A user account will be locked out if the user unsuccessfully tries to log in three consecutive times.

To unlock an account, follow these steps:

Step 1 From the WAAS Central Manager GUI, choose **Admin > AAA > Users**.

The User Accounts listing window appears and displays the status of each user account.

Note This window can only be accessed by users with administrator-level privileges.

Step 2 Click the **Edit** icon next to the user account that you want to modify.

The Modifying User Account window appears and displays a list of devices on which this account is locked out.

Step 3 Choose the device on which you want to unlock the account.

The list of device users appears.

Step 4 Choose the user or users to unlock, and click the **unlock** button.

Working with Passwords

The WAAS system features two levels of password policy: *standard* and *strong*. By default, the standard password policy is enabled.

To change the password policy, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Security > AAA > Password Policy Settings**.

Step 3 Check the **Enforce stringent password** check box to enable the strong password policy.

Step 4 In the **Maximum login retries** field, enter the maximum number of login attempts to be allowed before the user is locked out. The user remains locked out until cleared by the administrator. To clear a locked-out account, see the "Unlocking User Accounts" section on page 8-8.

Step 5 Click **Submit** to save your changes.

To configure password policy from the CLI, use the **authentication strict-password-policy** global configuration command.

When the standard password policy is enabled, user passwords must meet the following requirements:

- The password must be 1 to 31 characters long.
- The password can include both uppercase and lowercase letters (A–Z and a–z) and numbers (0–9).
- The password cannot contain the characters ' " | (apostrophe, double quote, or pipe) or any control characters.

When the strong password policy is enabled, user passwords must meet the following requirements:

- The password must be 8 to 31 characters long.
- The password can include both uppercase and lowercase letters (A–Z and a–z), numbers (0–9), and special characters including ~`!@#$%^&*()_+-=[]\{}|;:,</>.
- The password cannot contain the characters ' " | (apostrophe, double quote, or pipe) or any control characters.
- The password cannot contain all the same characters (for example, **99999**).
- The password cannot contain consecutive characters (for example, **12345**).
- The password cannot be the same as the username.
- Each new password must be different from the previous 12 passwords. User passwords expire within 90 days.
- The password cannot contain dictionary words.

A user account will be locked out after the configured number of failed login attempts (the default is three). The user remains locked-out until cleared by the administrator. To clear a locked-out account, see the "Unlocking User Accounts" section on page 8-8.

Working with Roles

The WAAS Central Manager GUI allows you to create roles for your WAAS system administrators so that each administrator can focus on configuring and managing a specific WAAS service. For example, you can set up a role that allows an administrator to create and modify application policies but does not allow the administrator to make any other changes to the system.

You can think of a role as a set of enabled services. Make sure you have a clear idea of the services that you want the role to be responsible for because you will select these services when you create the role. Once you create the role, you can assign the role to existing accounts as described later in this chapter.

A role can give read and write or read-only access to each enabled service.

Each user account or group can be assigned to zero or more roles. Roles are not inherited or embedded. The WAAS Central Manager provides one predefined role, known as the admin role. The admin role has access to all services, similar to a CLI user that has privilege level 15. Without the admin role, a user will not be able to perform all administrative tasks.

Note Assigning the admin role to a user does not change the user privilege level to 15. The user must also have privilege level 15 in order to perform administrative tasks.

Assigning the admin role to a user grants read and write permission to all Device Manager GUI pages.

WAAS can dynamically assign a role to users based on their membership in a group as defined on an external TACACS+ or Windows domain authentication server. To take advantage of this feature, you must define user group names on the WAAS Central Manager that match the user groups defined on the external authentication server and you must assign a role to the user groups on the WAAS Central Manager. For more information on user groups, see the "Working with User Groups" section on page 8-17.

Note For user groups authenticated on a TACACS+ server to gain access to the Device Manager GUI, the user group must be configured with the admin role and the user intending to access the Device Manager GUI must first log in to the Central Manager, which creates a member account on the Central Manager and the WAE. Periodically, member accounts of a user group are removed from the Central Manager database to reduce database load, so after a period (60 days by default) of no Central Manager activity, a user will need to log in again to the Central Manager before accessing the Device Manager GUI. The cdm.remoteuser.deletionDaysLimit system property controls the removal interval.

This section contains the following topics:

Creating a New Role

To create a new role, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > Roles**.

The Roles listing window appears.

Step 2 Click the **Create New Role** icon from the taskbar.

The Creating New Role window appears.

Step 3 In the Name field, enter the name of the role.

The name cannot contain characters other than letters, numbers, period, hyphen, underscore, and space.

Step 4 Check the check box next to the services that you want this role to manage.

The check boxes in this window are tri-state check boxes. When there is a check in the box, it means that the user will have read and write access to the listed service. Click the check box again to change the indicator to a square partially filling the check box. This indicator means that the user will have read-only access to the service. An empty square signifies no access to the service.

To expand the listing of services under a category, click the folder, and then check the check box next to the services that you want to enable for this role. To choose all the services under one category simultaneously, check the check box next to the top-level folder for those services.

Table 8-4 lists the services that you can enable for a role.

Table 8-4 Description of the WAAS Services

Service	Description
Home	Allows this role to view, configure, and manage the system dashboard and settings in the Configure, Monitor, and Admin menus of the WAAS Central Manager GUI in the Home (global) context. Under each folder you can select the subpages that you want this role to manage.
Device Groups	Allows this role to view, configure, and manage the settings and subpages for the various device groups in the WAAS Central Manager GUI in the device group context.
Devices	Allows this role to view, configure, and manage the settings and subpages for various kinds of devices in the WAAS Central Manager GUI in the device context.
AppNav Clusters	Allows this role to view, configure, and manage the settings and subpages in the WAAS Central Manager GUI in the AppNav Cluster context.
Locations	Allows this role to view, configure, and manage the settings and subpages in the WAAS Central Manager GUI in the Location context.
All Devices	Allows this role to access all the devices in your WAAS network. If this service is not enabled, the user account will only have access to the devices associated with the domain that you assign to the account. Selecting this service allows you to skip the following tasks when setting up a roles-based account: • Creating and maintaining a domain that contains all the devices in your network. • Assigning to the account the domain that contains all the devices.
All Device Groups	Allows this role to access all the device groups in your WAAS network. If this service is not enabled, then the user account will only have access to the device groups associated with the domain that you assigned to the account. Selecting this service allows you to skip the following tasks when setting up a roles-based account: • Creating and maintaining a domain that contains all the device groups in your network. • Assigning to the account the domain that contains all the device groups.
All AppNav Clusters	Allows this role to access all the AppNav Clusters in your WAAS network. If this service is not enabled, the user account will only have access to the AppNav Clusters associated with the domain that you assign to the account. Selecting this service allows you to skip the following tasks when setting up a roles-based account: • Creating and maintaining a domain that contains all the AppNav Clusters in your network. • Assigning to the account the domain that contains all the AppNav Clusters.

Table 8-4 Description of the WAAS Services (continued)

Service	Description
Monitoring API	Allows this role to access monitoring APIs through HTTPS requests. For more information, see the *Cisco Wide Area Application Services API Reference*.
System Status	Allows this role to access the device Alarms panel. For more information about device alarms, see Chapter 17, "Monitoring and Troubleshooting Your WAAS Network."

Step 5 (Optional) Enter any comments about this role in the Comments field.

Step 6 Click **Submit** to save your settings.

Assigning a Role to a User Account

After you create a role, you need to assign the role to an account (or a user group). If you create an account but do not assign a role to the account, that account can log into the WAAS Central Manager GUI but no data will be displayed and the configuration pages will not be available.

Note The admin user account, by default, is assigned to the role that allows access to all entities in the system. It is not possible to change the role for this user account.

To assign one or more roles to a user account group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > Users** (or **Admin > AAA > User Groups**).

The User Accounts (or User Groups) window appears with all configured user accounts listed.

Step 2 Click the **Edit** icon next to the user account or group for which you want to assign roles.

The Modifying User Account (or Modifying User Group) window appears.

Step 3 Click the **Role Management** tab.

The Role Management window appears with all configured role names listed.

Step 4 Click the **Assign** icon (blue cross mark) that appears next to the role name that you want to assign to the selected user account or group.

Step 5 Click the **Unassign** (green tick mark) next to the role name to unassign a previously assigned role.

Note Click the **Assign all Roles** icon in the taskbar to assign all roles in the current window to a user account or group. Alternatively, click the **Remove all Roles** icon to unassign all roles associated with a user account or group.

Step 6 Click **Submit**.

A green tick mark appears next to the assigned roles and a blue cross mark appears next to the unassigned roles. The roles assigned to this user account or group will be listed in the Roles section in the Modifying User Account (or Modifying User Group) window.

Modifying and Deleting Roles

Note The admin user account, by default, is allowed access to all services and cannot be modified.

To modify or delete a role, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > Roles**.

The Roles window appears.

Step 2 Click the **Edit** icon next to the name of the role you want to change or delete.

The Modifying Role window appears. You can modify the role as follows:

- To delete this role, click the **Delete** icon in the taskbar.
- To edit this role, make the necessary changes to the fields, and click **Submit**.
- To enable a service for this role, check the check box next to the services that you want. To disable a previously selected service, uncheck the check box next to the service you want to disable. To choose all the services under one category simultaneously, check the check box next to the top-level service.

Viewing Role Settings

You might want to view role settings before assigning a role to a particular user account or group.

To view role settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > Users** (or **Admin > AAA > User Groups**).

The User Accounts (or User Groups) window appears with all configured user accounts or groups listed.

Step 2 Click the **Edit** icon next to the user account or group that you want to view.

The Modifying User Account (or Modifying User Group) window appears.

Step 3 Click the **Role Management** tab.

The Role Management window appears.

Step 4 Click the **View** icon next to the role that you want to view.

The Viewing Role window appears, which displays the role name, comments about this role, and the services that are enabled for this role.

Step 5 After you have finished viewing the settings, click **Close**.

Working with Domains

A WAAS *domain* is a collection of device groups or WAEs that make up the WAAS network. A role defines which services a user can manage in the WAAS network, but a domain defines the device groups, WAEs, or file server dynamic shares that are accessible and configurable by the user.

Note A WAAS domain is not the same as a DNS domain or Windows domain.

When you create a domain, you choose the type of entities that can be associated with the domain. Entity types include Devices, Device Groups, or None (for file server dynamic shares). For file server dynamic shares, the dynamic shares are assigned in the dynamic shares configuration, as described in the "Creating Dynamic Shares for the CIFS Accelerator" section on page 12-9.

WAAS can dynamically assign a domain to a user based on their membership in a group as defined on an external TACACS+ or Windows domain authentication server. To take advantage of this feature, you must define user group names on the WAAS Central Manager that match the user groups defined on the external authentication server and you must assign a domain to the user groups on the WAAS Central Manager. For more information on user groups, see the "Working with User Groups" section on page 8-17.

This section contains the following topics:

Creating a New Domain

To create a new domain, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > Domains**.

The Domains listing window appears.

Step 2 Click the **Create New Domain** icon in the taskbar.

The Creating New Domain window appears.

Step 3 In the Name field, enter the name of the domain.

Step 4 From the Entity Type drop-down list, choose the entity type that you want to assign to the domain. Entity choices include Devices, Device Groups, and None. Choose **None** if this domain is used for a file server dynamic share.

Step 5 (Optional) In the Comments field, enter any comments about this domain.

Step 6 Click **Submit**.

If the entity type you chose has not already been assigned to the domain, then a message indicating that the entity type has not been assigned appears.

Step 7 Assign an entity to this domain as described in the section that follows, "Adding an Entity to a Domain". If you chose None for the Entity Type, do not assign an entity to the domain, instead, the entity is used in a dynamic share configuration, as described in the "Creating Dynamic Shares for the CIFS

Accelerator" section on page 12-9.

For a domain used in a dynamic share configuration, you must assign the domain to each user that needs to edit the dynamic share configuration, as described in the "Assigning a Domain to a User Account" section on page 8-15. Only users assigned to the domain will be able to edit the dynamic share.

Adding an Entity to a Domain

Once you have created a domain, you can assign an entity to the domain. An entity is either a collection of devices or a collection of device groups. You do not need to assign an entity to a domain that is used for a file server dynamic share, where the entity type is None.

To add an entity to a domain, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > Domains**.

Step 2 Click the **Edit** icon next to the domain that you want to modify.

Step 3 Click the **Entity Management** tab.

The *Entity_name* Assignments for Domain window for the current domain appears.

You can filter your view of the items in the list. Filtering enables you to find items in the list that match the criteria that you set.

You can add or remove entities from the domain as follows:

- To add an entity to the current domain, click the **Assign** icon (blue cross mark) next to the entity that you want to add. A green tick mark appears next to the selected entity when you submit the settings.

  Alternatively, to add all entities to the selected domain, click the **Assign all** icon in the taskbar.
- To remove an entity from the current domain, click the **Unassign** icon (green tick mark) next to the name of the entity that you want to remove from the domain. A blue cross mark appears next to the unassigned entity after you submit the settings.

  Alternatively, to remove all entities from the domain, click the **Remove all** icon in the taskbar.

Step 4 Click **Submit**.

Green check marks appear next to the entities that you assigned to the domain.

Step 5 Assign the domain to an account as described in the section that follows.

Assigning a Domain to a User Account

Assigning a domain to an account or user group specifies the entities (devices or device groups) or file server dynamic shares that the account or user group can access.

When working with a domain of type None that is used for dynamic file shares, you will need a user account for every user that needs to edit the dynamic share configuration. If you are using external authentication of users on TACACS+ or Windows domain servers, you can use user groups to more easily assign WAAS domains to users, see the "Working with User Groups" section on page 8-17.

Note If the role that you assigned to an account or group has the All Devices or All Device Groups service enabled, you do not need to assign a domain to the account or group. The account or group can automatically access all the devices and/or device groups in the WAAS system. For more information, see Table 8-4 on page 8-11.

To assign a domain to a user account or group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > Users** (or **Admin > AAA > User Groups**).

The User Accounts (or User Groups) window appears with all configured user accounts or groups listed.

Step 2 Click the **Edit** icon next to the user account or group for which you want to assign domains.

The Modifying User Account (or Modifying User Group) window appears.

Step 3 Click the **Domain Management** tab.

The Domain Management window appears with all configured domains and their entity types listed.

Step 4 Click the **Assign** icon (blue cross mark) that appears next to the domain name that you want to assign to the selected user account or group.

To dissociate an already associated domain from the user account or group, click the **Unassign** (green tick mark) next to the domain name.

Note To assign all domains in the current window to a user account or group, click the **Assign all Domains** icon in the taskbar. Alternatively, to unassign all domains associated with a user account or group, click the **Remove all Domains** icon.

Step 5 Click **Submit**.

A green check mark appears next to the assigned domains, and a blue cross mark appears next to the unassigned domains. The domains assigned to a user account or group are listed in the Domains section in the Modifying User Account (or Modifying User Group) window.

Modifying and Deleting Domains

To modify or delete an existing domain, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > Domains**.

The Domains window appears.

Step 2 Click the **Edit** icon next to the domain that you want to modify.

The Modifying Domain window appears. You can modify the domain as follows:

- To delete the domain, click the **Delete** icon in the taskbar and then click **OK** to confirm the deletion.
- To modify a domain, make the necessary changes to the fields and click **Submit**.

Viewing Domains

To view the domain configuration for a particular user account or group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > Users** (or **Admin > AAA > User Groups**).

The User Accounts (or User Groups) window appears with all configured user accounts or groups listed.

Step 2 Click the **Edit** icon next to the user account or group for which you want to view the domain configuration.

The Modifying User Account (or Modifying User Group) window appears.

Step 3 Click the **Domain Management** tab.

The Domain Management window appears.

Step 4 Click the **View** (eyeglass) icon next to the domain name to view details about the domain.

The Viewing Domain window appears and displays the domain name, entity type, comments about this domain, and entities assigned to this domain.

Step 5 After you have finished viewing the settings, click **Close**.

Working with User Groups

If you are using external authentication of users on TACACS+ or Windows domain servers (not RADIUS servers), you may want to create user groups. By creating user group names that match the user groups that you have defined on the external authentication server, WAAS can dynamically assign roles and WAAS domains to users based on their membership in a group as defined on the external authentication server. You do not need to define a role or WAAS domain for each user individually; instead, you define roles and WAAS domains for the user groups, and the user is assigned the roles and WAAS domains that are defined for the groups to which they belong.

Note The dynamic assignment of roles and WAAS domains based on external user groups requires a TACACS+ server that supports shell custom attributes. For example, these are supported in Cisco ACS 4.x and 5.1 and later.

WAAS reads group membership information for each user from the external authentication server.

This section contains the following topics:

Creating a New User Group

To create a new user group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > User Groups**.

The User Groups listing window appears.

Step 2 Click the **Create New User Groups** icon in the taskbar.

The Creating New User Group window appears.

Step 3 In the Name field, enter the name of the user group.

Ensure that the name matches the name of a user group defined on the external authentication server that you are using. Name matching is case sensitive.

Note A user group name cannot contain the following characters: # + " < > , (comma). A user group name cannot consist solely of numbers, periods (.), or spaces. Any leading periods, asterisks (*), or spaces are cropped.

Step 4 (Optional) In the Comments field, enter any comments about this user group.

Step 5 Click **Submit**.

Step 6 Assign a role or WAAS domain to this user group as described in the sections that follow.

Assigning Roles to a User Group

After you create a user group, you need to assign a role to the group. If you create a user group but do not assign a role to the group, the users in that group can log into the WAAS Central Manager GUI but no data will be displayed and the configuration pages will not be available.

To assign one or more roles to a user group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > User Groups**.

The User Groups window appears with all configured user groups listed.

Step 2 Click the **Edit** icon next to the user group for which you want to assign roles.

The Modifying User Group window appears.

Step 3 Click the **Role Management** tab.

The Role Management for User Group window appears with all configured role names listed.

Step 4 Click the **Assign** icon (blue cross mark) that appears next to the role name that you want to assign to the selected user group.

Step 5 Click the **Unassign** (green tick mark) next to the role name to unassign a previously assigned user group role.

Note Click the **Assign all Roles** icon in the taskbar to assign all roles in the current window to a user group. Alternatively, click the **Remove all Roles** icon to unassign all roles associated with a user group.

Step 6 Click **Submit**.

A green tick mark appears next to the assigned roles and a blue cross mark appears next to the unassigned roles. The roles assigned to this user group will be listed in the Roles section in the Modifying User Group window.

Assigning Domains to a User Group

Assigning a WAAS domain to a user group specifies the entities (devices or device groups) that the users who are members of the user group can manage.

Note If the role that you assigned to a user group has the All Devices or All Device Groups service enabled, you do not need to assign a domain to the user group. The users in that group can automatically access all the devices and/or device groups in the WAAS system. For more information, see Table 8-4 on page 8-11.

To assign a domain to a user group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA> User Groups**.

The User Groups window appears with all configured user groups listed.

Step 2 Click the **Edit** icon next to the user group for which you want to assign domains.

The Modifying User Group window appears.

Step 3 Choose the **Domain Management** tab.

The Domain Management for User Group window appears with all configured domains and their entity types listed.

Step 4 Click the **Assign** icon (blue cross mark) that appears next to the domain name that you want to assign to the selected user group.

To dissociate an already associated domain from the user group, click the **Unassign** (green tick mark) next to the domain name.

Note To assign all domains in the current window to a user group, click the **Assign all Domains** icon in the taskbar. Alternatively, to unassign all domains associated with a user group, click the **Remove all Domains** icon.

Step 5 Click **Submit**.

A green check mark appears next to the assigned domains, and a blue cross mark appears next to the unassigned domains. The domains assigned to a user group are listed in the Domains section in the Modifying User Group window.

Modifying and Deleting a User Group

To modify an existing user group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > AAA > User Groups**.

The User Groups window appears.

Step 2 Click the **Edit** icon next to the user group that you want to modify.

The Modifying User Group window appears. You can delete or edit user groups as follows:

Note This window can be accessed only by users with administrator-level privileges.

- To delete the user group, click the **Delete** icon in the taskbar, and then click **OK** to confirm the deletion.
- To edit the user group, make the necessary changes to the name and comment information, and click **Submit**.
- To change the Roles assigned to the user group, click the **Role Management** tab, make the necessary changes to the roles, and click **Submit**.
- To change the Domains assigned to the user group, click the **Domain Management** tab, make the necessary changes to the domains, and click **Submit**.

Viewing User Groups

To view all user groups, choose **Admin > AAA > User Groups** from the WAAS Central Manager GUI. The User Groups window displays all the user groups in the management database. From this window, you can also create groups as described in the "Creating a New User Group" section on page 8-18.

CHAPTER 9

Creating and Managing IP Access Control Lists for WAAS Devices

This chapter describes how to use the Wide Area Application Services (WAAS) Central Manager GUI to centrally create and manage Internet Protocol (IP) access control lists (ACLs) for your WAAS devices.

This chapter contains the following sections:

Note You must log in to the WAAS Central Manager GUI using an account with admin privileges to view, edit, or create IP ACL configurations.

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

About IP ACLs for WAAS Devices

In a centrally managed WAAS network environment, administrators need to be able to prevent unauthorized access to various devices and services. IP ACLs can filter packets by allowing you to permit or deny IP packets destined for a WAAS device.

The WAAS software supports standard and extended ACLs that allow you to restrict access to a WAAS device. The WAAS software can use the following types of ACLs:

- Interface ACL—Applied on the built-in, port channel, standby, and inline group interfaces. This type of ACL is intended to control management traffic (Telnet, SSH, and Central Manager GUI). The ACL rules apply only to traffic that is destined for the WAE or originates from the WAE, not WCCP transit traffic. Use the **ip access-group** interface configuration command to apply an interface ACL.
- Interception ACL—Applied globally to the WAAS device. This type of ACL defines what traffic is to be intercepted. Traffic that is permitted by the ACL is intercepted and traffic that is denied by the ACL is passed through the WAE. Use the **interception access-list** global configuration command to apply an interception ACL. For more information on using interception ACLs, see the "Configuring

Interception Access Control Lists" section on page 5-28.

- WCCP ACL—Applied on inbound WCCP redirected traffic to control access between an external server and external clients. The WAE is acting like a firewall. Use the **wccp access-list** global configuration command to apply a WCCP ACL.
- SNMP ACL—Applied on the SNMP agent to control access to the SNMP agent by an external SNMP server that is polling for SNMP MIBs or SNMP statistics. Use the **snmp-server access-list** global configuration command to apply an SNMP ACL.
- Transaction-logs flow ACL—Applied on the transaction logging facility to restrict the transactions to be logged. Use the **transaction-logs flow access-list** global configuration command to apply a transaction log ACL.

The following examples illustrate how interface ACLs can be used in environments that have WAAS devices:

- A WAAS device resides on the customer premises and is managed by a service provider, and the service provider wants to secure the device for its management only.
- A WAAS device is deployed anywhere within the enterprise. As with routers and switches, the administrator wants to limit access to Telnet, SSH, and the WAAS Central Manager GUI to the IT source subnets.

To use ACLs, you must first configure ACLs and then apply them to specific services or interfaces on the WAAS device. The following are some examples of how interface ACLs can be used in various enterprise deployments:

- An application layer proxy firewall with a hardened outside interface has no ports exposed. ("Hardened" means that the interface carefully restricts which ports are available for access primarily for security reasons. Because the interface is outside, many types of attacks are possible.) The WAAS device's outside address is globally accessible from the Internet, while its inside address is private. The inside interface has an ACL to limit Telnet, SSH, and GUI access.
- A WAE that is using WCCP is positioned on a subnet off the Internet router. Both the WAE and the router must have IP ACLs. IP access lists on routers have the highest priority followed by IP ACLs that are defined on the WAEs.

Note We strongly recommend that you use the WAAS Central Manager GUI instead of the WAAS CLI to centrally configure and apply ACLs to your WAAS devices. For more information, see the "Creating and Managing IP ACLs for WAAS Devices" section on page 9-2.

Creating and Managing IP ACLs for WAAS Devices

This section provides guidelines and an example of how to use the WAAS Central Manager GUI to create and manage IP ACLs for your WAAS devices.

When you create an IP ACL, you should note the following important points:

- IP ACL names must be unique within the device.
- IP ACL names must be limited to 30 characters and contain no white space or special characters.
- Each WAAS Central Manager device can manage up to 50 IP ACLs and a total of 500 conditions per device.

- When the IP ACL name is numeric, numbers 1 through 99 denote standard IP ACLs and numbers 100 through 199 denote extended IP ACLs. IP ACL names that begin with a number cannot contain nonnumeric characters.
- The WAAS Central Manager GUI allows the association of standard IP ACLs with SNMP and WCCP. Any device that attempts to access one of these applications associated with an ACL must be on the list of trusted devices to be allowed access.
- You can associate any previously configured standard IP ACL with SNMP and WCCP; however, you can associate an extended IP ACL only with the WCCP application.
- You can delete an IP ACL, including all conditions and associations with network interfaces and applications, or you can delete only the IP ACL conditions. Deleting all conditions allows you to change the IP ACL type if you choose to do so. The IP ACL entry continues to appear in the IP ACL listing; however, it is in effect nonexistent.
- If you specify an empty ACL for any of the ACL types used by WAAS, it has the effect of permitting all traffic.

To use the WAAS Central Manager GUI to create and modify an IP ACL for a single WAE, associate an IP ACL with an application, and then apply it to an interface on the WAE, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Network > TCP/IP Settings > IP ACL**.

The IP ACL window appears. By default, there are no IP ACLs defined for a WAE. The IP ACL window indicates if there are currently no IP ACLs configured for the WAE.

Step 3 In the taskbar, click the **Create a new IP ACL** icon.

The Creating New IP ACL window appears. Fill in the fields as follows:

- In the Name field, enter a name (for example, test1), observing the naming rules for IP ACLs.

  By default, this new IP ACL is created as a standard ACL.

  Note IP ACL names must be unique within the device, must be limited to 30 characters, and cannot contain any white spaces or special characters.

- If you want to change this default setting and create this new ACL as an extended ACL, choose **Extended** from the ACL Type drop-down list.

Step 4 Click **Submit** to save the IP ACL named test1. IP ACLs without any conditions defined do not appear on the individual devices.

Step 5 Add conditions to the standard IP ACL named test1 that you just created:

a. In the taskbar, click the **Create New Condition** icon.

The Creating New Condition window appears. (See Figure 9-1.)

Note The number of available fields for creating IP ACL conditions depends on the type of IP ACL that you have created, either standard or extended.

Figure 9-1 Creating a New Condition for an Extended IP ACL Window

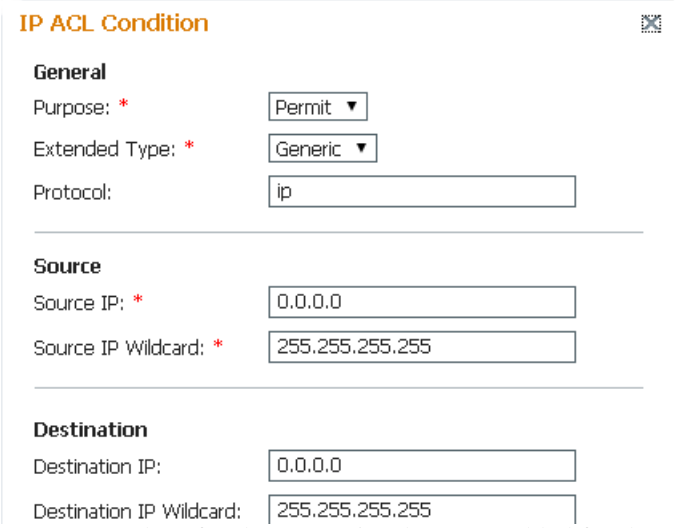

b. Enter values for the properties that are enabled for the type of IP ACL that you are creating, as follows:

- To set up conditions for a standard IP ACL, go to Step 6.
- To set up conditions for an extended IP ACL, go to Step 7.

Step 6 Set up conditions for a standard IP ACL:

a. From the drop-down list, choose a purpose (**Permit** or **Deny**).

b. In the Source IP field, enter the source IP address.

c. In the Source IP Wildcard field, enter a source IP wildcard address.

d. Click **Submit** to save the condition.

The Modifying IP ACL window reappears, displaying the condition and its configured parameters in tabular format.

e. To add another condition to the IP ACL, repeat the steps.

f. To reorder your list of conditions from the Modifying IP ACL window, use the Up or Down Arrows in the Move column, or click a column heading to sort by any configured parameter.

Note The order of the conditions listed in the WAAS Central Manager GUI becomes the order in which IP ACLs are applied to the device.

g. When you have finished adding conditions to the IP ACL, and you are satisfied with all your entries and the order in which the conditions are listed, click **Submit** in the Modifying IP ACL window to commit the IP ACL to the device database.

A green "Change submitted" indicator appears in the lower right corner of the Modifying IP ACL window to indicate that the IP ACL is being submitted to the device database. Table 9-1 describes the fields in a standard IP ACL.

Table 9-1 Standard IP ACL Conditions

Field	Default Value	Description
Purpose[1]	Permit	Specifies whether a packet is to be passed (**Permit**) or dropped (**Deny**).
Source IP[1]	0.0.0.0	Number of the network or host from which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Source IP Wildcard[1]	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.

1. Required field.

Step 7 Set up conditions for an extended IP ACL:

a. From the drop-down list, choose a purpose (**Permit** or **Deny**).

b. From the Extended Type drop-down list, choose **Generic**, **TCP**, **UDP**, or **ICMP**. (See Table 9-2.)

Table 9-2 Extended IP ACL Conditions

Field	Default Value	Description
Purpose[1]	Permit	Specifies whether a packet is to be passed or dropped. Choices are Permit or Deny.
Extended Type[1]	Generic	Specifies the Internet protocol to be applied to the condition. When selected, the GUI window refreshes with applicable field options enabled. The options are generic, TCP, UDP, or ICMP.

1. Required field.

After you choose a type of extended IP ACL, various options become available in the GUI, depending on what type you choose.

c. In the fields that are enabled for the chosen type, enter the data. (For more information, see Table 9-4 on page 9-8 through Table 9-7 on page 9-10.)

d. Click **Submit** to save the condition.

The Modifying IP ACL window reappears, displaying the condition and its configured parameters in tabular format.

e. To add another condition to the IP ACL, repeat the steps.

f. To reorder your list of conditions from the Modifying IP ACL window, use the Up or Down Arrows in the Move column, or click a column heading to sort by any configured parameter.

Note The order of the conditions listed in the WAAS Central Manager GUI becomes the order in which IP ACLs are applied to the device.

g. When you have finished adding conditions to the IP ACL, and you are satisfied with all your entries and the order in which the conditions are listed, click **Submit** in the Modifying IP ACL window to commit the IP ACL to the device database.

A green "Change submitted" indicator appears in the lower right corner of the Modifying IP ACL window to indicate that the IP ACL is being submitted to the device database.

Step 8 Modify or delete an individual condition from an IP ACL:

a. Click the **Edit** icon next to the name of the IP ACL that you want to modify. The Modifying IP ACL window appears, listing all the conditions that are currently applied to the IP ACL.

b. Click the **Edit Condition** icon next to the condition that you want to modify or delete. The Modifying Condition window appears.

c. To modify the condition, change any allowable field as necessary.

d. To delete the condition, click the **Trash** (**Delete IP ACL Condition**) icon in the taskbar.

e. To reorder your list of conditions, use the Up or Down arrows in the Move column, and click **Submit**.

Step 9 Associate a standard IP ACL with SNMP or WCCP:

a. Click the **Edit** icon next to the name of the device for which you want to associate a standard IP ACL with SNMP or WCCP.

b. Choose **Configure > Network > TCP/IP Settings > IP ACL Feature Usage**. The IP ACL Feature Settings window appears.

c. From the drop-down lists, choose the name of an IP ACL for SNMP or WCCP. (For more details, see Table 9-3.) If you do not want to associate an IP ACL with one of the applications, choose **Do Not Set**.

Table 9-3 IP ACL Feature Settings

WAAS Central Manager GUI Parameter	Function
SNMP	Associates a standard IP ACL with SNMP. This option is supported for WAAS devices that are operating as a WAE or a WAAS Central Manager device.
WCCP	Associates any IP ACL with WCCP Version 2. This option is only supported for WAAS devices that are operating as a WAE and not as a WAAS Central Manager device. WCCP is only supported on WAEs; it is not supported on a WAAS Central Manager device.

d. Click **Submit** to save the settings.

Step 10 Apply an IP ACL to an interface:

a. Click the **Edit** icon next to the name of the device for which you want to apply an IP ACL to an interface on the WAE.

b. Choose **Configure > Network > Network Interfaces**.

The Network Interfaces window for the device appears. This window displays all the interfaces available on that device.

c. Click the **Edit** icon next to the name of the interface to which you want to apply an IP ACL. The Network Interface settings window appears.

d. From the Inbound ACL drop-down list at the bottom of the window, choose the name of an IP ACL.

e. From the Outbound ACL drop-down list, choose the name of an ACL.

The only network interface properties that can be altered from the WAAS Central Manager GUI are the inbound and outbound IP ACLs. All other property values are populated from the device database and are read-only in the WAAS Central Manager GUI.

Step 11 Click **Submit** to save the settings.

Step 12 To use an IP ACL to define the traffic that should be intercepted, see the "Configuring Interception Access Control Lists" section on page 5-28.

Step 13 (Optional) Delete an IP ACL:

a. Click the **Edit** icon next to the name of the device that has the IP ACL that you want to delete.

b. Choose **Configure > Network > TCP/IP Settings > IP ACL**.

c. Click the **Edit** icon next to the name of the IP ACL that you want to delete (for example, test1).

The Modifying IP ACL window appears. If you created conditions for the IP ACL, you have two options for deletion:

- **Delete ACL**—Removes the IP ACL, including all conditions and associations with network interfaces and applications.
- **Delete All Conditions**—Removes all the conditions, while preserving the IP ACL name.

d. To delete the entire IP ACL, click the large **Trash** (**Delete ACL**) icon in the taskbar. You are prompted to confirm your action. Click **OK**. The record is deleted.

e. To delete only the conditions, click the small **Delete All Conditions** Trash/List icon in the taskbar. When you are prompted to confirm your action, click **OK**. The window refreshes, conditions are deleted, and the ACL Type field becomes available.

To define an IP ACL from the CLI, you can use the **ip access-list** global configuration command, and to apply the IP ACL to an interface on the WAAS device, you can use the **ip access-group** interface configuration command. To configure the use of an IP ACL for SNMP, you can use the **snmp-server access-list** global configuration command. To specify an IP ACL that the WAE applies to the inbound WCCP redirected traffic that it receives, you can use the **wccp access-list** global configuration command. To configure an interception ACL, you can use the **interception access-list** global configuration command.

List of Extended IP ACL Conditions

When you define a condition for an extended IP ACL, you can specify the Internet protocol to be applied to the condition (as described in Step 7 in the "Creating and Managing IP ACLs for WAAS Devices" section on page 9-2).

The list of extended IP ACL conditions are as follows:

- Generic (See Table 9-4.)
- TCP (See Table 9-5.)
- UDP (See Table 9-6.)
- ICMP (See Table 9-7.)

Table 9-4 Extended IP ACL Generic Condition

Field	Default Value	Description
Purpose[1]	Permit	Specifies whether a packet is to be passed (**Permit**) or dropped (**Deny**).
Extended Type[1]	Generic	Matches any Internet protocol.
Protocol	ip	Internet protocol (**gre**, **icmp**, **ip**, **tcp**, or **udp**). To match any Internet protocol, use the keyword **ip**.
Source IP[1]	0.0.0.0	Number of the network or host from which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Source IP Wildcard[1]	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.
Destination IP	0.0.0.0	Number of the network or host to which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Destination IP Wildcard	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.

1. Required field.

Table 9-5 Extended IP ACL TCP Condition

Field	Default Value	Description
Purpose[1]	Permit	Specifies whether a packet is to be passed (**Permit**) or dropped (**Deny**).
Extended Type[1]	TCP	Matches the TCP Internet protocol.
Established	Unchecked (false)	When checked, a match with the ACL condition occurs if the TCP datagram has the ACK or RST bits set, indicating an established connection. Initial TCP datagrams used to form a connection are not matched.
Source IP[1]	0.0.0.0	Number of the network or host from which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Source IP Wildcard[1]	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.
Source Port 1	0	Decimal number or name of a TCP port. Valid port numbers are 0 to 65535. Valid TCP port names are as follows: ftp, ftp-data, https, mms, netbios-dgm, netbios-ns, netbios-ss, nfs, rtsp, ssh, telnet, and www.

Table 9-5 Extended IP ACL TCP Condition (continued)

Field	Default Value	Description
Source Operator	range	Specifies how to compare the source ports against incoming packets. Choices are <, >, ==, !=, or range.
Source Port 2	65535	Decimal number or name of a TCP port. See Source Port 1.
Destination IP	0.0.0.0	Number of the network or host to which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Destination IP Wildcard	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.
Destination Port 1	0	Decimal number or name of a TCP port. Valid port numbers are 0 to 65535. Valid TCP port names are as follows: ftp, ftp-data, https, mms, netbios-dgm, netbios-ns, netbios-ss, nfs, rtsp, ssh, telnet, and www.
Destination Operator	range	Specifies how to compare the destination ports against incoming packets. Choices are <, >, ==, !=, or range.
Destination Port 2	65535	Decimal number or name of a TCP port. See Destination Port 1.

1. Required field.

Table 9-6 Extended IP ACL UDP Condition

Field	Default Value	Description
Purpose[1]	Permit	Specifies whether a packet is to be passed (**Permit**) or dropped (**Deny**).
Extended Type[1]	UDP	Matches the UDP Internet protocol.
Established	—	Not available for UDP.
Source IP[1]	0.0.0.0	Number of the network or host from which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Source IP Wildcard[1]	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.
Source Port 1	0	Decimal number or name of a UDP port. Valid port numbers are 0 to 65535. Valid UDP port names are as follows: bootpc, bootps, domain, mms, netbios-dgm, netbios-ns, netbios-ss, nfs, ntp, snmp, snmptrap, tacacs, tftp, and wccp.
Source Operator	range	Specifies how to compare the source ports against incoming packets. Choices are <, >, ==, !=, or range.

Table 9-6 Extended IP ACL UDP Condition (continued)

Field	Default Value	Description
Source Port 2	65535	Decimal number or name of a UDP port. See Source Port 1.
Destination IP	0.0.0.0	Number of the network or host to which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Destination IP Wildcard	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.
Destination Port 1	0	Decimal number or name of a UDP port. Valid port numbers are 0 to 65535. Valid UDP port names are as follows: bootpc, bootps, domain, mms, netbios-dgm, netbios-ns, netbios-ss, nfs, ntp, snmp, snmptrap, tacacs, tftp, and wccp.
Destination Operator	range	Specifies how to compare the destination ports against incoming packets. Choices are <, >, ==, !=, or range.
Destination Port 2	65535	Decimal number or name of a UDP port. See Destination Port 1.

1. Required field.

Table 9-7 Extended IP ACL ICMP Condition

Field	Default Value	Description
Purpose[1]	Permit	Specifies whether a packet is to be passed (**Permit**) or dropped (**Deny**).
Extended Type[1]	ICMP	Matches the ICMP Internet protocol.
Source IP[1]	0.0.0.0	Number of the network or host from which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Source IP Wildcard[1]	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.
Destination IP	0.0.0.0	Number of the network or host to which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Destination IP Wildcard	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.

Table 9-7 Extended IP ACL ICMP Condition (continued)

Field	Default Value	Description
ICMP Param Type[1]	None	Choices are **None**, **Type/Code**, or **Msg**. **None**—Disables the ICMP Type, Code, and Message fields. **Type/Code**—Allows ICMP messages to be filtered by ICMP message type and code. Also enables the ability to set an ICMP message code number. **Msg**—Allows a combination of type and code to be specified using a keyword. Activates the ICMP message drop-down list. Disables the ICMP Type field.
ICMP Message[1]	administratively-prohibited	Allows a combination of ICMP type and code to be specified using a keyword chosen from the drop-down list.
ICMP Type[1]	0	Number from 0 to 255. This field is enabled when you choose **Type/Code**.
Use ICMP Code[1]	Unchecked	When checked, enables the ICMP Code field.
ICMP Code[1]	0	Number from 0 to 255. Message code option that allows ICMP messages of a particular type to be further filtered by an ICMP message code.

1. Required field.

CHAPTER 10

Configuring Other System Settings

This chapter describes how to perform other system tasks such as setting the system clock, modifying the default system configuration settings, and enabling alarm overload detection, after you have done a basic configuration of your WAAS device. This chapter also describes how to register and manage WAAS Express devices.

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

This chapter contains the following sections:

Modifying Device Properties

The WAAS Central Manager GUI allows you to make the following changes to the properties of a WAE device:

- Rename the device
- Assign a new location to the device
- Assign an IP address to be used for management traffic to the device

- Deactivate or activate the device

You can also use the WAAS Central Manager GUI to check the status of a device to determine if it is online, pending, or inactive.

You can only rename a WAAS Central Manager device from the GUI.

To modify a device's properties, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose *device-name* > **Activation**.

The Device Activation window appears with fields for editing the properties of the selected device.

For a WAAS Central Manager device, the only fields that you can change in this window are the name and NetBIOS name of the device. In addition, the device IP address and role are displayed.

Step 3 Under the General Configuration heading, set or modify the following device properties:

- To change the hostname of the device, enter a new name in the Name field. This name must conform to the following rules:
  - The name must use only alphanumeric characters and hyphens (-).
  - The first and last character must be a letter or a digit.
  - Maximum length is 30 characters.
  - Names are case insensitive.
  - The following characters are considered illegal and cannot be used when naming a device: @, #, $,%, ^, &, *, (), |, \""/, <>.
- To activate or deactivate the device, check or uncheck the **Activate** check box.When this box is checked, the device is activated for centralized management through the WAAS Central Manager GUI.

  You can also click the **Deactivate** icon in the task bar to deactivate the device. Deactivating a device allows you to replace the device in the event of a hardware failure without losing all of its configuration settings.
- To change the NetBIOS name of the device, enter the new NetBIOS name for the device in the provided field. The NetBIOS name must not consist of only numbers; it must include some letters. This field is not displayed for WAAS Express devices.

Step 4 Under the Locality heading, set or change the location by choosing a new location from the **Location** drop-down list. To create a location for this device, see the "Creating Locations" section on page 3-10.

Step 5 Under the Management Interface Configuration with NAT heading, configure the NAT settings using the following fields:

- Check the **Use WAE's primary IP Address** check box to enable the WAAS Central Manager to use the IP address configured on the primary interface of the device to communicate with devices in the WAAS network that are behind a NAT firewall. This check box is not displayed for WAAS Express devices.
- Allow the WAAS Central Manager to communicate with devices in the WAAS network that are behind the NAT firewall using an explicitly configured IP address, by entering the IP address of the device in the Management IP field. You also need to enter this address in scenarios where the primary interface for a WAE is set to an inline group interface and management traffic is configured on a separate IP address (either on a secondary IP address on the same inline group interface or on a built-in interface).

- In the Port field, enter the port number for the management IP address. If the HTTPS server configured on a WAAS Express device is using a different port than the default of 443, configure the same port here.

Note If the WAAS Central Manager cannot contact a device using the primary IP address, it attempts to communicate using the Management IP address.

Step 6 In the Comments field, enter any comments that you want to appear for this device.

Step 7 Click **Submit**.

Managing Software Licenses

WAAS software version 4.1.1 introduces software licenses that enable specific WAAS optimization and acceleration features. A software license must be installed and configured before the features that it enables will operate.

Table 10-1 lists the software licenses that may be purchased and the features that each license enables.

Table 10-1 WAAS Software Licenses

License	Description
Transport	Enables basic DRE, TFO, and LZ optimization. Cannot be configured if the Enterprise license is configured.
Enterprise	Enables the EPM, HTTP, MAPI, NFS, SSL, CIFS, SMB, ICA, and Windows Print application accelerators, the WAAS Central Manager, and basic DRE, TFO, and LZ optimization. Cannot be configured if the Transport license is configured.
Video	Enables the video application accelerator. Requires the Enterprise license to be configured first.
Virtual-Blade	Enables the virtualization feature. Requires the Enterprise license to be configured first.

Licenses are installed and managed only on individual WAE devices, not device groups. Not all licenses are supported on all devices. A WAAS Central Manager device requires only the Enterprise license and no other licenses can be configured.

Note WAAS Express licenses are managed by using the router CLI command **license install**, not from the WAAS Central Manager. WAAS Express devices do not use the same kind of licenses as WAAS devices do. They use a single license that enables the WAAS Express optimization feature.

To add a license to a WAE from the WAAS Central Manager, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*. (Do not choose a Central Manager device because you must use the CLI to manage licenses on Central Managers.)

Step 2 Choose **Admin > History > License Management**.

Step 3 Check the check box next to each license that you want to add.

Step 4 Click **Submit**.

To add licenses from the CLI, you can use the **license add** EXEC command.

To remove licenses from the CLI, you can use the **clear license** EXEC command.

To display the status of all licenses from the CLI, you can use the **show license** EXEC command.

The setup utility also configures licenses when you first set up a new WAAS device.

Enabling the Inetd RCP and FTP Services

Remote Copy Protocol (RCP) lets you download, upload, and copy configuration files between remote hosts and a switch. Unlike TFTP, which uses User Datagram Protocol (UDP), a connectionless protocol, RCP uses TCP, which is connection oriented. Inetd (an Internet daemon) is a program that listens for connection requests or messages for certain ports and starts server programs to perform the services associated with those ports. RCP copies files between devices.

RCP is a subset of the UNIX rshell service, which allows UNIX users to execute shell commands on remote UNIX systems. It is a UNIX built-in service. This service uses TCP as the transport protocol and listens for requests on TCP port 514. RCP service can be enabled on WAAS devices that use WAAS software.

To enable RCP and FTP services on a WAAS device, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Network > Network Services**. The Network Services window appears.

Step 3 Check the **Enable Rcp Service** check box to enable Inetd RCP services. By default, this option is disabled.

Note The Inetd daemon listens for FTP, RCP, and TFTP services. For Inetd to listen to RCP requests, it must be explicitly enabled for RCP service.

Step 4 Check the **Enable FTP Service** check box to enable the Inetd FTP service. By default, this option is disabled.

Step 5 Click **Submit** to save your changes.

A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking the Reset button. The Reset button is visible only when you have applied default or group settings to change the current device settings but you have not yet submitted the changes.

If you try to leave this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box only appears if you are using the Internet Explorer browser.

Configuring Date and Time Settings

This section explains how to configure date and time settings for your WAAS network devices and contains the following topics:

Configuring NTP Settings

The WAAS Central Manager GUI allows you to configure the time and date settings using a Network Time Protocol (NTP) host on your network. NTP allows the synchronization of time and date settings for the different geographical locations of the devices in your WAAS network, which is important for proper system operation and monitoring. On each WAAS device, be sure to set up an NTP server to keep the clocks synchronized.

To configure NTP settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Date/Time > NTP**. The NTP Settings window appears.

Step 3 Check the **Enable** check box to enable NTP settings. By default, this option is disabled.

Step 4 In the NTP Server field, enter a hostname or IP address.

Step 5 Click **Submit**.

Note Unexpected time changes can result in unexpected system behavior. We recommend reloading the system after configuring an NTP server or changing the system clock.

Configuring Time Zone Settings

If you have an outside source on your network that provides time services (such as a Network Time Protocol [NTP] server), you do not need to set the system clock manually. When manually setting the clock, enter the local time.

Note Two clocks exist in the system: the software clock and the hardware clock. The software uses the software clock. The hardware clock is used only at startup to initialize the software clock.

To configure the time zone on a device or device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Date/Time > Time Zone**. The Time Zone Settings window appears.

Step 3 To configure a standard time zone, follow these steps:

a. Under the Time Zone Settings section, click the **Standard Time Zone** radio button. The default is UTC (offset = 0) with no summer time configured. When you configure a standard time zone, the system is automatically adjusted for the UTC offset, and the UTC offset need not be specified.

   The standard convention for time zones uses a *Location/Area* format in which *Location* is a continent or a geographic region of the world and *Area* is a time zone region within that location.

b. From the drop-down list, choose a location for the time zone. (For an explanation of the abbreviations in this list, see Table 10-2.)

   The window refreshes, displaying all area time zones for the chosen location in the second drop-down list.

c. Choose an area for the time zone. The UTC offset is automatically set for standard time zones.

   Summer time is built-in for some standard time zones (mostly time zones within the United States), and will result an automatic change in the UTC offset during summer time. For a list of standard time zones that can be configured and their UTC offsets, see Table 10-3.

Step 4 To configure a customized time zone on the device, follow these steps:

a. Under the Time Zone Settings section, click the **Customized Time Zone** radio button.

b. In the Customized Time Zone field, specify the name of the time zone. The time zone entry is case-sensitive and can contain up to 40 characters including spaces. If you specify any of the standard time zone names, an error message is displayed when you click **Submit**.

c. For UTC Offset, choose the **+** or **–** sign from the first drop-down list to specify whether the configured time zone is ahead or behind UTC. Also, choose the number of hours (0–23) and minutes (0–59) offset from UTC for the customized time zone. The range for the UTC offset is from –23:59 to 23:59, and the default is 0:0.

Step 5 To configure customized summer time, follow these steps under the Customized Summer Time Savings section.

Note You can specify a customized summer time for both standard and customized time zones.

a. To configure absolute summer time, click the **Absolute Dates** radio button.

   You can configure a start date and end date for summer time in absolute dates or recurring dates. Absolute date settings apply only once and must be set every year. Recurring dates apply repeatedly for many years.

b. In the Start Date and End Date fields, specify the month (January through December), day (1–31), and year (1993–2032) on which summer time must start and end in mm/dd/yyyy format. Make sure that the end date is always later than the start date.

   Alternatively, click the **Calendar** icon next to the Start Date and End Date fields to display the Date Time Picker popup window. By default the current date is highlighted in yellow. In the Date Time Picker popup window, use the left or right arrow icons to choose the previous or following years, if required. Choose a month from the drop-down list. Click a day of the month. The chosen date is highlighted in blue. Click **Apply**. Alternatively, click **Set Today** to revert to the current day. The chosen date will be displayed in the Start Date and End Date fields.

c. To configure recurring summer time, click the **Recurring Dates** radio button.

d. From the Start Day drop-down list, choose a day of the week (**Monday**-**Sunday**) to start.

e. From the Start Week drop-down list, choose an option (**first**, **2nd**, **3rd**, or **last**) to set the starting week. For example, choose **first** to configure summer time to recur beginning the first week of the month or **last** to configure summer time to recur beginning the last week of the month.

f. From the Start Month drop-down list, choose a month (**January–December**) to start.

g. From the End Day drop-down list, choose a day of the week (**Monday–Sunday**) to end.

h. From the End Week drop-down list, choose an option (**first**, **2nd**, **3rd**, or **last**) to set the ending week. For example, choose **first** to configure summer time to end beginning the first week of the month or **last** to configure summer time to stop beginning the last week of the month.

i. From the End Month drop-down list, choose a month (**January–December**) to end.

Step 6 From the Start Time drop-down lists, choose the hour (0–23) and minute (0–59) at which daylight saving time should start. From the End Time drop-down lists, choose the hour (0–23) and minute (0–59) at which daylight saving time should end.

Start Time and End Time fields for summer time are the times of the day when the clock is changed to reflect summer time. By default, both start and end times are set at 00:00.

Step 7 In the Offset field, specify the minutes offset from UTC (0–1439). (See Table 10-3.)

The summer time offset specifies the number of minutes that the system clock moves forward at the specified start time and backward at the end time.

Step 8 Click the **No Customized Summer Time Configured** radio button to not specify a summer or daylight saving time for the corresponding time zone.

Step 9 Click **Submit** to save the settings.

A "Click Submit to Save" message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking the Reset button. The Reset button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

If you attempt to leave this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box only appears if you are using the Internet Explorer browser.

Table 10-2 Timezone Location Abbreviations

Time Zone	Expansion
CET	Central European Time
CST6CDT	Central Standard/Daylight Time
EET	Eastern European Time
EST	Eastern Standard Time
EST5EDT	Eastern Standard/Daylight Time
GB	Great Britain
GB-Eire	Great Britain/Ireland
GMT	Greenwich Mean Time
HST	Hawaiian Standard Time
MET	Middle European Time
MST	Mountain Standard Time
MST7MDT	Mountain Standard/Daylight Time
NZ	New Zealand
NZ-CHAT	New Zealand, Chatham Islands

Table 10-2 Timezone Location Abbreviations (continued)

Time Zone	Expansion
PRC	People's Republic of China
PST8PDT	Pacific Standard/Daylight Time
ROC	Republic of China
ROK	Republic of Korea
UCT	Coordinated Universal Time
UTC	Coordinated Universal Time
WET	Western European Time
W-SU	Middle European Time

Table 10-3 Timezone—Offset from UTC

Time Zone	Offset from UTC (in hours)
Africa/Algiers	+1
Africa/Cairo	+2
Africa/Casablanca	0
Africa/Harare	+2
Africa/Johannesburg	+2
Africa/Nairobi	+3
America/Buenos_Aires	–3
America/Caracas	–4
America/Mexico_City	–6
America/Lima	–5
America/Santiago	–4
Atlantic/Azores	–1
Atlantic/Cape_Verde	–1
Asia/Almaty	+6
Asia/Baghdad	+3
Asia/Baku	+4
Asia/Bangkok	+7
Asia/Colombo	+6
Asia/Dacca	+6
Asia/Hong_Kong	+8
Asia/Irkutsk	+8
Asia/Jerusalem	+2
Asia/Kabul	+4.30
Asia/Karachi	+5
Asia/Katmandu	+5.45
Asia/Krasnoyarsk	+7
Asia/Magadan	+11
Asia/Muscat	+4

Table 10-3 Timezone—Offset from UTC (continued)

Time Zone	Offset from UTC (in hours)
Asia/New Delhi	+5.30
Asia/Rangoon	+6.30
Asia/Riyadh	+3
Asia/Seoul	+9
Asia/Singapore	+8
Asia/Taipei	+8
Asia/Tehran	+3.30
Asia/Vladivostok	+10
Asia/Yekaterinburg	+5
Asia/Yakutsk	+9
Australia/Adelaide	+9.30
Australia/Brisbane	+10
Australia/Darwin	+9.30
Australia/Hobart	+10
Australia/Perth	+8
Australia/Sydney	+10
Canada/Atlantic	–4
Canada/Newfoundland	–3.30
Canada/Saskatchewan	–6
Europe/Athens	+2
Europe/Berlin	+1
Europe/Bucharest	+2
Europe/Helsinki	+2
Europe/London	0
Europe/Moscow	+3
Europe/Paris	+1
Europe/Prague	+1
Europe/Warsaw	+1
Japan	+9
Pacific/Auckland	+12
Pacific/Fiji	+12
Pacific/Guam	+10
Pacific/Kwajalein	–12
Pacific/Samoa	–11
US/Alaska	–9
US/Central	–6
US/Eastern	–5
US/East–Indiana	–5
US/Hawaii	–10

Table 10-3 Timezone—Offset from UTC (continued)

Time Zone	Offset from UTC (in hours)
US/Mountain	–7
US/Pacific	–8

UTC was formerly known as Greenwich Mean Time (GMT). The offset time (number of hours ahead or behind UTC) as displayed in the table is in effect during winter time. During summer time or daylight saving time, the offset may be different from the values in the table and is calculated and displayed accordingly by the system clock.

Configuring Secure Store Settings

Secure store encryption provides strong encryption and key management for your WAAS system. The WAAS Central Manager and WAE devices use secure store encryption for handling passwords, managing encryption keys, and for data encryption.

This section contains the following topics:

Secure Store Overview

With secure store encryption on the Central Manager or a WAE device, the WAAS system uses strong encryption algorithms and key management policies to protect certain data on the system. This data includes encryption keys used by applications in the WAAS system, CIFS accelerator passwords for prepositioning, user login passwords, NAM credentials, and certificate key files.

Secure store encryption on the Central Manager is always enabled and uses a password that is auto-generated or user-provided. This password is used to generate the *key encryption key* according to secure standards. The WAAS system uses the key encryption key to encrypt and store other keys generated on the Central Manager or WAE devices. These other keys are used for WAAS functions including disk encryption, SSL acceleration, or to encrypt and store CIFS accelerator credentials, and user passwords.

Data on the Central Manager is encrypted using a 256-bit key encryption key generated from the password and using SHA1 hashing and an AES 256-bit algorithm. When secure store is enabled on a WAE device the data is encrypted using a 256-bit key encryption key generated using SecureRandom, a cryptographically strong pseudorandom number generator.

Secure store encryption on a Central Manager uses one of the following modes:

- Auto-generated passphrase mode—The passphrase is automatically generated by the Central Manager and used to open the secure store after each system reboot. This is the default mode for new Central Manager devices or after the system has been reinstalled.
- User-provided passphrase mode—The passphrase is supplied by the user and must be entered after each system reboot to open the secure store. You can switch to this mode, and systems upgraded from versions prior to 4.4.1, with secure store initialized, are configured in this mode after upgrading to 4.4.1 or later.

To implement secure store your system must meet the following requirements:

- You must have a Central Manager configured for use in your network.
- Your WAE devices must be registered with the Central Manager.
- Your WAE devices must be online (have an active connection) with the Central Manager. This requirement applies only if you are enabling secure store on WAE devices.
- All Central Managers and WAE devices must be running WAAS software version 4.0.19 or higher.

To implement strong store encryption, follow these steps:

Step 1 Enable strong storage encryption on your primary Central Manager. See Enabling Secure Store Encryption on the Central Manager.

Step 2 Enable strong storage encryption on any standby Central Managers. See Enabling Secure Store Encryption on a Standby Central Manager.

Step 3 Enable strong storage encryption on WAE devices or WAE device groups. See Enabling Secure Store Encryption on a WAE Device. (Secure store must be enabled on the Central Manager before you enable it on the WAE devices.)

You can enable secure store independently on the Central Manager and on the WAE devices. To ensure full protection of your encrypted data, enable secure store on both the Central Manager and the WAE devices. You must enable secure store on the Central Manager first.

Note When you reboot the Central Manager, if secure store is in user-provided passphrase mode, you must manually open secure store encryption. All services that use the secure store (disk encryption, CIFS prepositioning, SSL acceleration, AAA, and so on) on the remote WAE devices do not operate properly until you enter the secure store password on the Central Manager to open secure store encryption.

Note the following considerations regarding the secure store:

- Passwords stored in the Central Manager database are encrypted using strong encryption techniques.
- CIFS prepositioning credentials are encrypted using the strong encryption key on the Central Manager and the WAE devices.
- Certificate key files are encrypted using the strong encryption key on the Central Manager.
- If a primary Central Manager fails, secure store key management is handled by the standby Central Manager. (Secure store mode must be enabled manually on the standby Central Manager.)
- Backup scripts back up the secure store passphrase mode (user-provided or auto-generated) of the device at the time of backup. Backup and restore are supported only on the Central Manager.

- If you have a backup made when the secure store was in user-provided passphrase mode and you restore it to a system where the secure store is in auto-generated passphrase mode, you must enter the user passphrase to proceed with the restore. After the restore, the system is in user-provided passphrase mode. If you have a backup made when the secure store was in auto-generated passphrase mode and you restore it to a system where the secure store is in user-provided passphrase mode, you do not need to enter a password. After the restore, the system is in auto-generated passphrase mode.
- When you enable secure store on a WAE device, the system initializes and retrieves a new encryption key from the Central Manager. The WAE uses this key to encrypt data such as CIFS prepositioning credentials and information on the disk (if disk encryption is also enabled).
- When you reboot the WAE after enabling secure store, the WAE retrieves the key from the Central Manager automatically, allowing normal access to the data that is stored in WAAS persistent storage. If key retrieval fails, a critical alarm is raised and secure store should be reopened manually. Until secure store is reopened, the WAE rejects configuration updates from the Central Manager if the updates contain CIFS preposition, dynamic share, or user configuration. Also, the WAE does not include preposition configuration in the updates that it sends to the Central Manager.
- While secure store encrypts certain system information, it does not encrypt the data on the hard drives. To protect the data disks, you must enable disk encryption separately. See the "Enabling Disk Encryption" section on page 16-30.

Enabling Secure Store Encryption on the Central Manager

Secure store is enabled by default on a new Central Manager, with a system-generated password that opens the secure store after the system boots. You do not need to do anything to enable secure store.

If a Central Manager is configured in user-provided passphrase mode, you must manually open the secure store after the system boots. To open secure store encryption on the Central Manager, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > Secure Store**. The Configure CM Secure Store window appears.

Step 2 Enter the secure store passphrase in the Current passphrase field under Open Secure Store.

Step 3 Click the **Open** button.

The secure store is opened. Data is encrypted using the key derived from the password.

To open the secure store from the CLI, use the **cms secure-store open** EXEC command.

Note Whenever you reboot a Central Manager that is configured in user-provided passphrase mode, you must reopen the secure store manually. All services that use the secure store (disk encryption, CIFS prepositioning, SSL acceleration, AAA, and so on) on the remote WAE devices do not operate properly until you enter the secure store password on the Central Manager to reopen the secure store. Switch to auto-generated passphrase mode to avoid having to reopen the secure store after each reboot.

Note When you enable secure store on the primary Central Manager in user-provided passphrase mode, you should enable secure store on the standby Central Manager as well. See Enabling Secure Store Encryption on a Standby Central Manager, page 10-13.

You can check the status of secure store encryption by entering the **show cms secure-store** command.

Enabling Secure Store Encryption on a Standby Central Manager

Note A standby Central Manager provides limited encryption key management support. If the primary Central Manager fails, the standby Central Manager provides only encryption key retrieval to the WAE devices but does not provide new encryption key initialization. Do not enable disk encryption or secure store on WAE devices when the primary Central Manager is not available.

The secure store passphrase mode on the primary Central Manager is replicated to the standby Central Manager (within the standard replication time). If the primary Central Manager is switched to auto-generated passphrase mode, the standby Central Manager secure store changes to the open state. If the primary Central Manager is switched to user-provided passphrase mode or the passphrase is changed, the standby Central Manager secure store changes to the initialized but not open state and an alarm is raised. You must manually open the secure store on the standby Central Manager.

To enable secure store encryption on a standby Central Manager when the primary Central Manager is in user-provided passphrase mode, open the secure store on the primary Central Manager and then use the CLI to execute the **cms secure-store open** EXEC mode command on the standby Central Manager:

Step 1 Enable secure store encryption on the primary Central Manager. See the "Enabling Secure Store Encryption on the Central Manager" section on page 10-12.

Step 2 Wait until the standby Central Manager replicates the data from the primary Central Manager.

The replication should occur in 60 seconds (default) or as configured for your system.

Step 3 Enter the **cms secure-store open** command on the standby Central Manager to activate secure store encryption.

The standby Central Manager responds with the "please enter pass phrase" message.

Step 4 Type the password and press **Enter**.

The standby Central Manager encrypts the data using secure store encryption.

Note Repeat Steps 3 and 4 for each standby Central Manager on your system.

You can check the status of secure store encryption by entering the **show cms secure-store** command.

Enabling Secure Store Encryption on a WAE Device

To enable secure store encryption on a WAE device, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Note The secure store status must be the same for all WAE devices in a device group. Either all WAE devices in the group must have secure store enabled, or all must have secure store disabled. Before you add a WAE device to a device group, set its secure store status to match the others. See the "Working with Device Groups" section on page 3-2.

Step 2 Choose **Configure > Security > Secure Store**. The Secure Store Settings window appears

Step 3 Check the **Initialize CMS Secure Store** box. (The Open CMS Secure Store box will be checked automatically.)

Step 4 Click **Submit** to activate secure store encryption.

A new encryption key is initialized on the Central Manager, and the WAE encrypts the data using secure store encryption.

To enable secure store from the CLI, use the **cms secure-store init** EXEC command.

Note If you have made any other CLI configuration changes on a WAE within the datafeed poll rate time interval (5 minutes by default) before executing the **cms secure-store** command, those prior configuration changes are lost and you must redo them.

Note When you enable or disable secure store on a device group, the changes do not take effect on all WAE devices simultaneously. When you view the WAE devices be sure to give the Central Manager enough time to update the status of each WAE device.

Changing Secure Store Passphrase Mode

The secure store can operate either in user-provided or auto-generated passphrase mode and you can switch between these modes.

To change from user-provided to auto-generated passphrase mode, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > Secure Store**.

Step 2 In the Switch to CM auto-generated passphrase mode area, enter the password in the Current passphrase field.

Step 3 Click the **Switch** button.

Step 4 Click **OK** in the confirmation message that appears.

The secure store is changed to auto-generated passphrase mode and remains in the open state.

To change from auto-generated to user-provided passphrase mode, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > Secure Store**.

Step 2 In the Switch to User-provided passphrase mode area, enter a password in the New passphrase field and reenter the password in the Confirm passphrase field.

The password must conform to the following rules:

- Be 8 to 64 characters in length
- Contain characters only from the allowed set: A-Za-z0-9~%'!#$^&*()|;:,"<>/
- Contain at least one digit
- Contain at least one lowercase and one uppercase letter

Step 3 Click the **Switch** button.

Step 4 Click **OK** in the confirmation message that appears.

The secure store is changed to user-provided passphrase mode and remains in the open state. If you have a standby Central Manager, you must manually open its secure store (see the "Enabling Secure Store Encryption on a Standby Central Manager" section on page 10-13).

To change secure store passphrase mode from the CLI, use the **cms secure-store mode** EXEC command.

Note Whenever you reboot a Central Manager that is configured in user-provided passphrase mode, you must reopen the secure store manually. All services that use the secure store (disk encryption, CIFS prepositioning, SSL acceleration, AAA, and so on) on the remote WAE devices do not operate properly until you enter the secure store password on the Central Manager to reopen the secure store. Switch to auto-generated passphrase mode to avoid having to reopen the secure store after each reboot.

Changing the Secure Store Encryption Key and Password

The secure store encryption password is used by the Central Manager to generate the encryption key for the encrypted data. If the Central Manager is configured for user-provided passphrase mode, you can change the password.

To change the password and generate a new encryption key on the Central Manager, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > Secure Store**.

Step 2 In the Change Secure Store passphrase area, in the Current passphrase field, enter the current password.

Step 3 In the New passphrase field, enter the new password.

The password must conform to the following rules:

- Be 8 to 64 characters in length
- Contain characters only from the allowed set: A-Za-z0-9~%'!#$^&*()|;:,"<>/
- Contain at least one digit
- Contain at least one lowercase and one uppercase letter

Step 4 In the Confirm passphrase field, enter the new password again.

Step 5 Click the **Change** button.

The WAAS device reencrypts the stored data using a new encryption key derived from the new password.

To change the password and generate a new encryption key on the Central Manager from the CLI, use the **cms secure-store change** EXEC command.

To generate a new encryption key for a WAE device from the WAAS Central Manager, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Security > Secure Store**.

Step 3 Check the **Change CMS Secure Store** box and then click **Submit**.

A new encryption key is generated in the Central Manager. The Central Manager replaces the encryption key in the WAE with the new key. The WAE re-encrypts the stored data using the new encryption key.

To configure the secure store encryption key from the CLI, use the **cms secure-store change** EXEC command.

Resetting Secure Store Encryption on a Central Manager

You can reset the secure store if you reload the Central Manager and you cannot open the secure store because it is configured in user-provided passphrase mode and you forget the secure store password. This procedure deletes all encrypted data, certificate and key files, and key manager keys. The secure store is reinitialized, configured in auto-generated passphrase mode, and opened.

To reset secure store encryption on a Central Manager, follow these steps:

Step 1 At the primary Central Manager CLI, enter the **cms secure-store reset** command to reset secure store encryption.

Step 2 Wait until the standby Central Manager replicates the data from the primary Central Manager.

The replication should occur in 60 seconds (default) or as configured for your system.

Step 3 Enter the **cms secure-store reset** command on the standby Central Manager if secure store is in the initialized and open state.

Step 4 From the primary Central Manager, reset all user account passwords, CIFS credentials, and NAM credentials.

For information on resetting user passwords, see the "Changing the Password for Another Account" section on page 8-7. For information on resetting dynamic share passwords, see the "Creating Dynamic Shares for the CIFS Accelerator" section on page 12-9. For information on resetting preposition passwords, see the "Creating a New Preposition Directive" section on page 12-12. For information on resetting NAM credentials, see the "Configuring the Basic Setup" section on page 15-3.

Step 5 On each WAE registered to the Central Manager, follow these steps:

a. If secure store is initialized and open, from the Central Manager, clear secure store (see the "Disabling Secure Store Encryption on a WAE Device" section on page 10-17). Or, from the CLI, enter the **cms secure-store clear** EXEC command.

b. From the Central Manager, initialize secure store (see the "Enabling Secure Store Encryption on a WAE Device" section on page 10-13) or from the CLI, enter the **cms secure-store init** EXEC command. (This step is needed only if you performed step 5a.)

c. Enter the **crypto pki managed-store initialize** command and restart the SSL accelerator.

d. If disk encryption is enabled, from the Central Manager, disable disk encryption (see the "Enabling Disk Encryption" section on page 16-30) or from the CLI, enter the **no disk encrypt enable** global configuration command.

e. If disk encryption had been enabled before step 5d, reload the device. After the reload, reenable disk encryption and reload the device again.

Note If the WAE is reloaded before doing Step 5, disk encryption, SSL acceleration, and secure store does not function properly. In this case, you must restore the WAE to factory defaults.

Step 6 From the primary Central Manager, reimport all certificate and key files for all the accelerated and peering services which are configured on the WAEs.

Disabling Secure Store Encryption on a WAE Device

To disable secure store encryption on a WAE device, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Security > Secure Store**. The Secure Store Settings window appears.

Step 3 Check the **Clear CMS Secure Store** box and then click Submit to disable secure store encryption and return to standard encryption.

You can also enter the **cms secure-store clear** command to disable secure store encryption and return to standard encryption.

To disable secure store on a WAE from the CLI, use the **cms secure-store clear** EXEC command.

Note Secure store cannot be disabled on a Central Manager.

Modifying the Default System Configuration Properties

The WAAS software comes with preconfigured system properties that you can modify to alter the default behavior of the system.

Table 10-4 describes the system configuration properties that you can modify.

Table 10-4 Descriptions for System Configuration Properties

System Property	Description
cdm.remoteuser.deletionDaysLimit	Maximum number of days since their last login after which external users will be deleted from the WAAS Central Manager database. For example, if cdm.remoteuser.deletionDaysLimit is set to 5, external users will be deleted from the database if the difference between their last login time and the current time is more than 5 days. The default is 60 days. External users are users that are defined in an external AAA server and not in the WAAS Central Manager. Any reports scheduled by such users are also deleted when the users are deleted.
cdm.session.timeout	Timeout in minutes of a WAAS Central Manager GUI session. The default is 10 minutes. If the session is idle for this length of time, the user is automatically logged out.
DeviceGroup.overlap	Status of whether a device can belong to more than one device group. The default is true (devices can belong to more than one device group).
System.datafeed.pollRate	Poll rate between a WAAS (or WAAS Express) device and the WAAS Central Manager (in seconds). The default is 300 seconds.
System.device.recovery.key	Device identity recovery key. This property enables a device to be replaced by another node in the WAAS network.
System.guiServer.fqdn	Scheme to use (IP address or FQDN) to launch the Device Manager GUI.
System.healthmonitor.collectRate	Collect and send rate in seconds for the CMS device health (or status) monitor. If the rate is set to 0, the health monitor is disabled. The default is 120 seconds.
System.lcm.enable	Local and central management feature (enable or disable). This property allows settings that are configured using the local device CLI or the WAAS Central Manager GUI to be stored as part of the WAAS network configuration data. The default is true. If this property is set to false (disabled), configuration changes made on a local device will not be communicated to the Central Manager and configurations done in the Central Manager will overwrite local device configurations. This setting applies to both WAAS and WAAS Express devices.
System.monitoring.collectRate	Rate at which a WAE collects and sends the monitoring report to the WAAS Central Manager (in seconds). For a WAAS Express device, this is the rate at which the Central Manager collects the monitoring data from the WAAS Express device. The default is 300 seconds (5 minutes). Reducing this interval impacts the performance of the WAAS Central Manager device.
System.monitoring.dailyConsolidationHour	Hour at which the WAAS Central Manager consolidates hourly and daily monitoring records. The default is 1 (1:00 a.m.).
System.monitoring.enable	WAAS and WAAS Express statistics monitoring (enable or disable). The default is true.
System.monitoring.maxDevicePerLocation	Maximum number of devices for which monitoring is supported in location level reports. The default is 25.

Table 10-4 Descriptions for System Configuration Properties (continued)

System Property	Description
System.monitoring.maxReports	Maximum number of completed or failed report instances to store for each custom report. The default is 10 report instances.
System.monitoring.monthlyConsolidationFrequency	How often (in days) the WAAS Central Manager consolidates daily monitoring records into monthly records. If this setting is set to 1, the WAAS Central Manager checks if consolidation needs to occur every day, but only performs consolidation if there is enough data for consolidation to occur. The default is 14 days. When a monthly data record is created, the corresponding daily records are removed from the database. Consolidation occurs only if there is at least two calendar months of data plus the consolidation frequency days of data. This ensures that the WAAS Central Manager always maintains daily data records for the past month and can display data on a day level granularity for the last week. For example, if data collection starts on February 2nd, 2006 and System.monitoring.monthlyConsolidationFrequencyis set to 14, then the WAAS Central Manager checks if there is data for the past two calendar months on the following days: Feb 16th, March 2nd, March 16th, and March 30th. No consolidation will occur because there is not enough data on these days. On April 13th, however, two calendar months of data exists. The WAAS Central Manager then consolidates data for the month of February and deletes the daily data records for that month.
System.monitoring.recordLimitDays	Maximum number of days of monitoring data to maintain in the system. The default is 1825 days.
System.monitoring.timeFrameSettings	Default time frame to be used for plotting all the charts. Settings saved by the user will not be changed. The default is Last Hour.
System.registration.autoActivation	Status of the automatic activation feature, which automatically activates WAAS and WAAS Express devices that are registered to the Central Manager. The default is true (devices are automatically registered).
System.rpc.timeout.syncGuiOperation	Timeout in seconds for the GUI synchronization operations for the Central Manager to WAE connection. The default is 50 seconds.
System.security.maxSimultaneousLogins	Maximum number of concurrent WAAS Central Manager sessions permitted for a user. Specify 0 (zero, the default) for unlimited concurrent sessions. A user must log off the Central Manager to end a session. If a user closes the browser without logging off, the session is not closed until after it times out after 120 minutes (the timeout is not configurable). If the number of concurrent sessions permitted also is exceeded for that user, there is no way for that user to regain access to the Central Manager GUI until after the timeout expires. This setting does not affect CLI access to the Central Manager device.
System.security.webApplicationFilter	Status of the web application filter, which rejects any javascript, SQL, or restricted special characters in input. The default is false.

Table 10-4 Descriptions for System Configuration Properties (continued)

System Property	Description
System.standby.replication.maxCount	Maximum number of statistics data records (in thousands) that will be replicated to a standby Central Manager. The range is 10 to 300. The default is 200 (200,000 records). We do not recommend increasing this number.
System.standby.replicationTimeout	Maximum number of seconds to wait for replication to a standby Central Manager. The range is 300 to 3600 seconds. The default is 900 seconds. We do not recommend decreasing this timeout.

To view or modify the value of a system property, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Configure > Global > System Properties**. The Config Properties window appears.

Step 2 Click the **Edit** icon next to the system property that you want to change. The Modifying Config Property window appears.

Step 3 From a drop-down list, enter a new value or choose a new parameter, depending on the system property that you want to change.

Step 4 Click **Submit** to save the settings

Configuring the Web Application Filter

Web Application Filter is a security feature that protects the WAAS Central Manager GUI against Cross-Site Scripting (XSS) attacks. XSS security issues can occur when an application sends data that originates from a user to a web browser without first validating or encoding the content, which can allow malicious scripting to be executed in the client browser, potentially compromising database integrity.

This security feature verifies that all application parameters sent from WAAS users are validated and/or encoded before populating any HTML pages.

This section contains the following topics:

Enabling the Web Application Filter

To enable the Web Application Filter, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Configure > Global > System Properties**. The Config Properties window appears.

Note You cannot enable this feature using the CLI. This feature is disabled by default.

Step 2 Click the Edit icon next to the system.security.webApplicationFilter entry.

The Modifying Config Property window appears.

Step 3 Choose **true** from the Value drop-down list to enable this feature.

A confirmation message appears to advise Central Manager and Device Manager users to log out and then back in after enabling this feature.

Step 4 Click **OK** and then **Submit**.

Step 5 Log out and then back in again.

Security Verification

The Web Application Filter feature verifies security using two methods, input verification and sanitization. Input validation validates all input data before accepting data. Sanitization prevents malicious configuration and scripts already present in the data from getting executed.

This section contains the following topics:

- Input Validation, page 10-21
- Sanitization, page 10-21

Input Validation

Input validation scans all data that is input to the Central/Device Manager database and is only configurable by the admin user.

Any input submitted using the Central Manager GUI that is suspicious of XSS is blocked. Blocked input results in a warning.

Input data is checked against the following XSS filter rules:

- Input is rejected if it contains a semicolon (;)
- Input is rejected if it is enclosed in angle brackets (<>)
- Input is rejected if it can be indirectly used to generate the above tags (&#60, &#62, %3c, %3e)

Sanitization

The sanitizer prevents malicious configuration and scripts from getting executed in the browser when there is an XSS attack on the database. Sanitization is not configurable by the user.

Configuration data coming from the Central Manager that is suspect for XSS is shown in red on the Device Groups > All Device Groups page.

Configuring Faster Detection of Offline WAAS Devices

You can detect offline WAAS devices more quickly if you enable the fast detection of offline devices. A WAAS device is declared as offline when it has failed to contact the WAAS Central Manager for a getUpdate (get configuration poll) request for at least two polling periods. (See the "About Faster Detection of Offline Devices" section on page 10-22 for more information about this feature.)

To configure fast detection of offline WAAS devices, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Configure > Global > Fast Device Offline Detection**. The Configure Fast Offline Detection window appears.

Note The fast detection of offline devices feature is in effect only when the WAAS Central Manager receives the first UDP heartbeat packet and a getUpdate request from a device.

Step 2 Check the **Enable** check box to enable the WAAS Central Manager to detect the offline status of devices quickly.

Step 3 In the Heartbeat Rate (Seconds) field, specify how often devices should transmit a UDP heartbeat packet to the WAAS Central Manager. The default is 30 seconds.

Step 4 In the Heartbeat Fail Count field, specify the number of UDP heartbeat packets that can be dropped during transmission from devices to the WAAS Central Manager before a device is declared offline. The default is 1.

Step 5 In the Heartbeat UDP Port field, specify the port number using which devices will send UDP heartbeat packets to the primary WAAS Central Manager. The default is port 2000.

The Maximum Offline Detection Time field displays the product of the failed heartbeat count and heartbeat rate.

Maximum Offline Detection Time = Failed heartbeat count * Heartbeat rate

If you have not enabled the fast detection of offline devices feature, then the WAAS Central Manager waits for at least two polling periods to be contacted by the device for a getUpdate request before declaring the device to be offline. However, if you enable the fast detection of offline devices feature, then the WAAS Central Manager waits until the value displayed in the Maximum Offline Detection Time field is exceeded.

If the WAAS Central Manager receives the Cisco Discovery Protocol (CDP) from a device, then the WAAS Central Manager GUI displays the device as offline after a time period of 2* (heartbeat rate) * (failed heartbeat count).

Step 6 Click **Submit**.

Note Any changes to the Configure Fast WAE offline detection page in the Central Manager could result in devices temporarily appearing to be offline. Once the configuration changes are propagated to the devices, they show as online again.

About Faster Detection of Offline Devices

Communication between the WAAS device and WAAS Central Manager using User Datagram Protocol (UDP) allows faster detection of devices that have gone offline. UDP heartbeat packets are sent at a specified interval from each device to the primary WAAS Central Manager in a WAAS network. The primary WAAS Central Manager tracks the last time that it received a UDP heartbeat packet from each device. If the WAAS Central Manager has not received the specified number of UDP packets, it displays

the status of the nonresponsive devices as offline. Because UDP heartbeats require less processing than a getUpdate request, they can be transmitted more frequently, and the WAAS Central Manager can detect offline devices much faster.

You can enable or disable this feature, specify the interval between two UDP packets, and configure the failed heartbeat count. Heartbeat packet rate is defined as the interval between two UDP packets. Using the specified heartbeat packet rate and failed heartbeat count values, the WAAS Central Manager GUI displays the resulting offline detection time as a product of heartbeat rate and failed heartbeat count. If the fast detection of offline devices is enabled, the WAAS Central Manager detects devices that are in network segments that do not support UDP and uses getUpdate (get configuration poll) request to detect offline devices.

By default, the feature to detect offline devices more quickly is not enabled.

Configuring Alarm Overload Detection

WAAS devices can track the rate of incoming alarms from the Node Health Manager. If the rate of incoming alarms exceeds the high-water mark (HWM), then the WAAS device enters an alarm overload state. This situation occurs when multiple applications raise alarms at the same time to report error conditions. When a WAAS device is in an alarm overload state, the following occurs:

- SNMP traps for subsequent alarm raise and clear operations are suspended. The trap for the raise alarm-overload alarm and the clear alarm-overload alarm are sent; however, traps related to alarm operations between the raise alarm-overload alarm and the clear alarm-overload alarm operations are suspended.
- Alarm overload raise and clear notifications are not blocked. The alarm overload state is communicated to SNMP and the Configuration Management System (CMS). However, in the alarm overload state, SNMP and the CMS are not notified of individual alarms. The information is only available by using the CLI.
- The WAAS device remains in an alarm overload state until the rate of incoming alarms decreases to the point that the alarm rate is less than the low-water mark (LWM).
- If the incoming alarm rate falls below the LWM, the WAAS device comes out of the alarm overload state and begins to report the alarm counts to SNMP and the CMS.

When the WAAS device is in an alarm overload state, the Node Health Manager continues to record the alarms being raised on the WAAS device and keeps a track of the incoming alarm rate. Alarms that have been raised on a WAAS device can be listed using the **show alarm** CLI commands that are described in the *Cisco Wide Area Application Services Command Reference*.

To configure alarm overload detection for a WAAS device (or device group), follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Monitoring > Alarm Overload Detection**. The Alarm Overload Detection Settings window appears.

Step 3 Uncheck the **Enable Alarm Overload Detection** check box if you do not want to configure the WAAS device (or device group) to suspend alarm raise and clear operations when multiple applications report error conditions. This check box is checked by default.

Step 4 In the Alarm Overload Low Water Mark (Clear) field, enter the number of incoming alarms per second below which the WAAS device comes out of the alarm overload state.

The low-water mark is the level up to which the number of alarms must drop before alarms can be restarted. The default value is 1. The low-water mark value should be less than the high-water mark value.

Step 5 In the Alarm Overload High Water Mark (Raise) field, enter the number of incoming alarms per second above which the WAAS device enters the alarm overload state. The default value is 10.

Step 6 Click **Submit** to save the settings.

To configure alarm overload detection from the CLI, you can use the **alarm overload-detect** global configuration command.

Configuring the E-mail Notification Server

You can schedule reports to be generated periodically, and when they are generated, a link to the report can be e-mailed to one or more recipients. (For details, see the "Managing Reports" section on page 17-43.)

To enable e-mail notification, you must configure e-mail server settings for the WAAS Central Manager by following these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*. You must choose a Central Manager device.

Step 2 Choose **Configure > Monitoring > Email Notification**. The Configure Email Server Details window appears.

Step 3 In the Mail Server Hostname field, enter the hostname of the SMTP e-mail server that is to be used to send e-mail.

Note Only SMTP mail servers are supported. If any other type of mail server is configured, the email notification fails.

Step 4 In the Mail Server Port field, enter the port number. The default is port 25.

Step 5 In the Server Username field, enter a valid e-mail account username.

Step 6 In the Server Password field, enter the password for the e-mail account.

Step 7 In the From Address field, enter the e-mail address shown as the sender of the e-mail notification.

Step 8 Click **Submit**.

Using IPMI over LAN

Intelligent Platform Management Interface (IPMI) over LAN provides remote platform management service for WAVE-294/594/694/7541/7571/8541 appliances. IPMI is an open standard technology that defines how administrators monitor system hardware and sensors, control system components, and retrieve logs of important system events to conduct remote management and recovery. IPMI runs on the Baseboard Management Controller (BMC) and operates independently of WAAS. After IPMI over LAN

is set up and enabled on WAAS, authorized users can access BMC remotely even when WAAS becomes unresponsive or the device is powered down but connected to a power source. You can use an IPMI v2 compliant management utility, such as ipmitool or OSA SMbridge, to connect to the BMC remotely to perform IPMI operations.

The IPMI over LAN feature provides the following remote platform management services:

- Supports the power on, power off, and power cycle of the WAAS appliance.
- Monitors the health of the WAAS hardware components by examining Field Replaceable Unit (FRU) information and reading sensor values.
- Retrieves logs of important system events to conduct remote management and recovery.
- Provides serial console access to the WAAS appliance over the IPMI session.
- Support for IPMI Serial over LAN (SoL)—IPMI SoL enables a remote user to access a WAAS appliance through a serial console through an IPMI session.

IPMI over LAN and IPMI SoL features can be configured using CLI commands and include the following:

- Configuring IPMI LAN interface
- Configuring IPMI LAN users
- Configuring security settings for remote IPMI access
- Enabling/disabling IPMI over LAN
- Enabling/disabling IPMI SoL
- Restoring the default settings for the BMC LAN channel
- Displaying the current IPMI over LAN and IPMI SoL configurations

For more information on configuring IPMI over LAN, see the "Configuring BMC for Remote Platform Management" section on page 10-26.

BMC Firmware Update

IPMI over LAN requires that a specific BMC firmware version be installed on the device. The minimum supported BMC firmware versions are:

- WAVE-294/594/694—48a
- WAVE-7541/7571/8541—26a

WAAS appliances shipped from the factory with WAAS version 4.4.5 or later do have the correct firmware installed. If you are updating a device that was shipped with an earlier version of WAAS software, you must update the BMC firmware, unless it was updated previously.

To determine if you are running the correct firmware version, use the **show bmc info** command. The following example displays the latest BMC firmware version installed on the device (48a here):

```
wave# show bmc info
Device ID                 : 32
Device Revision           : 1
Firmware Revision         : 0.48                    <<<<< version 48
IPMI Version              : 2.0
Manufacturer ID           : 5771
Manufacturer Name         : Unknown (0x168B)
Product ID                : 160 (0x00a0)
Product Name              : Unknown (0xA0)
Device Available          : yes
Provides Device SDRs      : no
Additional Device Support :
```

```
    Sensor Device
    SDR Repository Device
    SEL Device
    FRU Inventory Device
Aux Firmware Rev Info     :
    0x0b
    0x0c
    0x08
    0x0a                                          <<<<< a
. . .
```

If a BMC firmware update is needed, you can download it from cisco.com at the Wide Area Application Service (WAAS) Firmware download page (registered customers only). The firmware binary image is named waas-bmc-installer-48a-48a-26a-k9.bin.

You can use the following command to update the firmware from the image file that is available through FTP on your network:

copy ftp install *ip-address remotefiledir* **waas-bmc-installer-48a-48a-26a-k9.bin**

The update process automatically checks the health status of the BMC firmware. If BMC firmware corruption is detected, BMC is recovered during the BMC firmware update procedure. The complete update process can take several minutes and the device may appear unresponsive but do not interrupt the process or power cycle the device. After the update is complete, you must reload the device.

After the device reboots, you can verify the firmware version by using the **show bmc info** command.

BMC recovery and BMC firmware update restores the factory defaults on the BMC and all the current IPMI over LAN configurations are erased.

If BMC firmware corruption happens, a critical alarm is raised.

Configuring BMC for Remote Platform Management

This section describes the minimum steps needed to enable IPMI over LAN and IPMI SoL to conduct remote platform management. This section includes the following topics:

- Enabling IPMI Over LAN
- Enabling IPMI SoL

Enabling IPMI Over LAN

To enable IPMI over LAN, perform the following steps using the **bmc lan** command:

Step 1 Change the default BMC LAN IP address.

Step 2 Change the password for the BMC default user, which is user 2.

Step 3 Enable IPMI over LAN.

Step 4 Access the BMC from a remote client over IPMI session v2.0 using the username and password for the number 2 user. The default cipher suite used to access the BMC is 3, which specifies RAKP-HMAC-SHA1 authentication, HMAC-SHA1-96 integrity, and AES-CBC-128 encryption algorithms.

Step 5 To access the BMC over a IPMI session v1.5, change the user 2 IPMI-session-version setting from v2.0 to v1.5.

Enabling IPMI SoL

To enable IPMI SoL, perform the following steps:

Step 1 On the WAAS device, configure and enable IPMI over Lan (IoL).

Step 2 On the remote client make sure that the BMC user can do IoL operations successfully over IPMI session v2.0.

Step 3 On the remote client, change the baud-rate of the terminal to match the WAAS console baud rate of 9600 bps.

Step 4 On the WAAS device, enable IPMI SoL.

Step 5 On the remote client, if the IPMI management tool is ipmitool, check the SoL payload status of the specific BMC user with the following command:
ipmitool -I lanplus -H *bmc-ip-address* **-U** *bmc-user-name* **sol payload status 1** *bmc-user-userid*

For example:

```
# ipmitool -I lanplus -H 2.1.4.70 -U user3 sol payload status 1 3
Password:
User 3 on channel 1 is disabled
```

Step 6 If the SoL payload is disabled for this user, enable the SoL payload for this user with the following command:
ipmitool -I lanplus -H *bmc-ip-address* **-U** *bmc-user-name* **sol payload enable 1** *bmc-user-userid*

For example:

```
# ipmitool -I lanplus -H 2.1.4.70 -U user3 sol payload enable 1 3
Password:
# ipmitool -I lanplus -H 2.1.4.70 -U user3 sol payload status 1 3
Password:
User 3 on channel 1 is enabled
```

Step 7 On the remote client, use the following command to open the serial console to the WAAS device:
ipmitool -I lanplus -H *bmc-ip-address* **-U** *bmc-user-name* **sol activate**

Step 8 On the remote client, you have now entered the console session of the WAAS device. When you are done, use the ~. escape character to terminate the connection.

Managing WAAS Express Devices

You can use the WAAS Central Manager to manage WAAS Express devices, which are Cisco ISR G2 routers deployed with the WAAS Express software. The WAAS Express software implements a subset of the WAAS appliance functionality, providing basic optimization and HTTP accelerator express, CIFS accelerator express, and SSL accelerator express. The Central Manager menu displays a subset of the full menu when a WAAS Express device is selected as the context.

The Central Manager and a WAAS Express device communicate using the HTTPS protocol. To establish communication between a WAAS Central Manager and a WAAS Express device, you must register the WAAS Express device with the Central Manager. Using the Central Manager GUI to register a WAAS Express device is a more simplified method than using the CLI.

- Registering a WAAS Express Device Using the GUI, page 10-28
- Registering a WAAS Express Device Using the CLI, page 10-29

Registering a WAAS Express Device Using the GUI

To register a WAAS Express device, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > Security > WAAS Express > Registration**. The WAAS Express Registration window appears. (See Figure 10-1.)

Note To register a WAAS Express device using the Central Manager GUI, SSH must be enabled on the WAAS Express device.

Figure 10-1 WAAS Express Registration

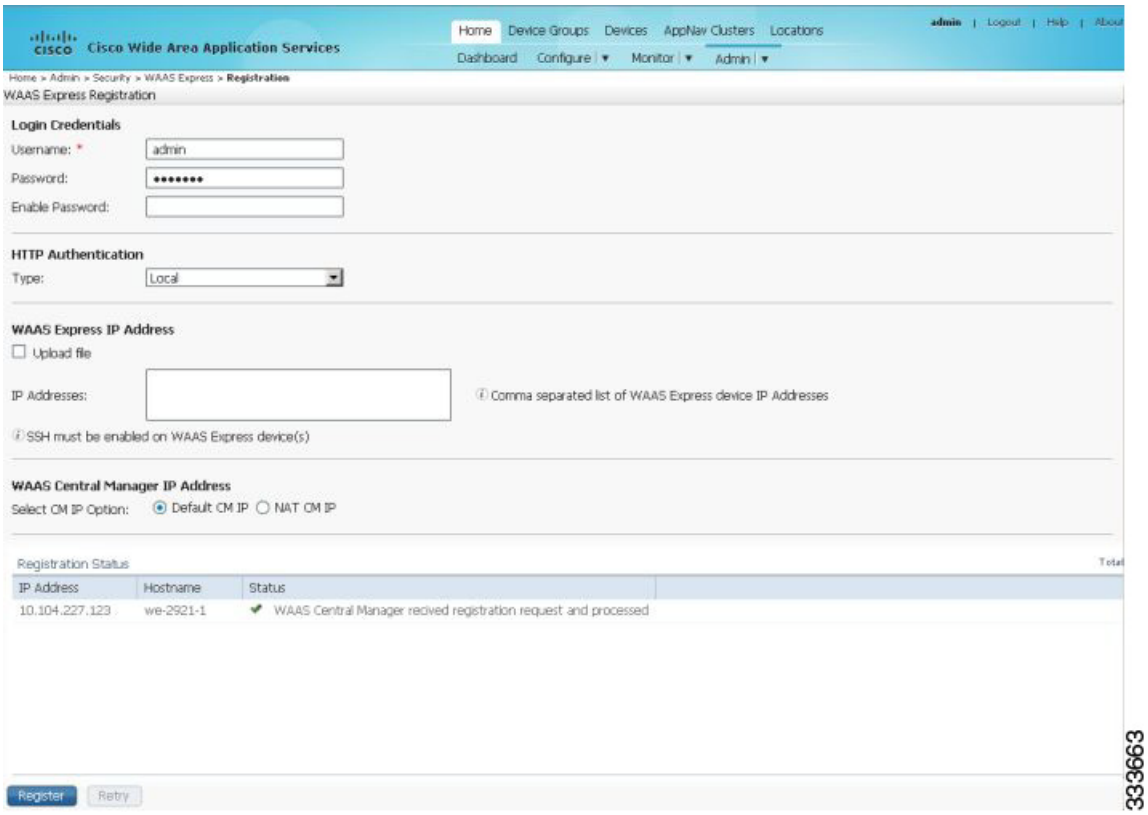

Step 2 Configure the login credentials by entering the username, password, and enable password.

Step 3 Enter the HTTP Authentication, local or AAA.

Step 4 Enter the WAAS Express IP addresses to register. The IP address, hostname, and status are displayed in the Registration Status table.

You may also upload a CSV file that contains a list of IP addresses to register. To upload a list, check the Upload file check box and either browse to the file or enter the filename. Each IP address must be on a separate line.

Step 5 Select the Central Manager IP address option, either default or NAT.

Step 6 Click the Register button and verify that the registration status was successful.

You may view the results in the log file: /local/local1/errlog/waasx-audit.log

Step 7 The final step is to install a permanent WAAS software license. This function is not supported using the Central Manager GUI. You must obtain and copy the WAAS license to a location accessible to the **license** command on the WAAS Express device.

The following is an example of the license command used on the WAAS Express device to install the WAAS license:

```
waas-express#license install ftp://infra/licenses/FHH122500AZ_20100811190225615.lic
```

This example uses FTP to get and install the license but there are various options available for this command. Choose one that best suits your deployment.

Registering a WAAS Express Device Using the CLI

You can register the WAAS Express device with the Central Manager using the CLI by completing the steps outlined in Table 10-5.

Table 10-5 Checklist for Registering a WAAS Express Device using the CLI

Task	Additional Information and Instructions
1. Configure a username and password.	The same username and password are configured on the WAAS Express device and the Central Manager, so the Central Manager can log in to the WAAS Express device for management purposes. For more information, see the "Configuring a User" section on page 10-30.
2. Import the primary Central Manager administrative server certificate into the WAAS Express device.	The WAAS Express device requires the Central Manager certificate for secure HTTPS server communication. For more information, see the "Importing the Central Manager Certificate" section on page 10-31.
3. Configure a WAAS Express device certificate.	The Central Manager device requests this WAAS Express certificate for secure HTTPS server communication. For more information, see the "Configuring a WAAS Express Device Certificate" section on page 10-32.
4. Enable the secure HTTP server with user authentication.	Enables the Central Manager and WAAS Express device to communicate. For more information, see the "Enabling the HTTP Secure Server on the WAAS Express Device" section on page 10-32.
5. Install a permanent WAAS software license.	Allows the WAAS Express software to operate on the router. For more information, see the "Installing a License on the WAAS Express Device" section on page 10-33.
6. Configure an NTP server.	Keeps the time synchronized between the WAAS Express device and the Central Manager. For more information, see the "Configuring an NTP Server" section on page 10-33.
7. Register the WAAS Express device with the Central Manager.	Registers the WAAS Express device with the Central Manager. For more information, see the "Registering the WAAS Express Device" section on page 10-34.

The following sections describe these steps in detail.

Configuring a User

The first step in setting up your WAAS Express device and Central Manager to communicate is to configure the same user on the WAAS Express device and the Central Manager.

To configure a user, follow these steps:

Step 1 Log in to the WAAS Express device CLI.

Step 2 Configure a local user with privilege level 15 on the WAAS Express device by using the **username** IOS configuration command:

```
waas-express#config t
Enter configuration commands, one per line. End with CNTL/Z.
waas-express(config)#username cisco privilege 15 password 0 cisco
waas-express(config)#exit
```

Alternatively, you can configure an external TACACS+ or RADIUS user; see details after this procedure.

Step 3 Save the running configuration:

```
waas-express#write memory
Building configuration...
[OK]
```

Step 4 From the WAAS Central Manager menu, choose **Admin > Security > WAAS Express > Global Credentials**. The WAAS Express Global Credentials window appears.

On the Central Manager, you can define global WAAS Express credentials that apply to all WAAS Express devices, or you can define credentials at the device group or individual device level. This procedure shows how to configure global credentials. To configure device group or individual device credentials, you must first complete the WAAS Express registration process and then configure this setting for a WAAS Express device group or device. Device and device group credentials have precedence over global credentials.

Step 5 In the Username field, enter the same username that you defined on the WAAS Express device.

Note The username field is optional if you are not using local or AAA authentication for the HTTP server on the WAAS Express device; that is, if you use the default HTTP server configuration of **ip http authentication enable**. (See the "Enabling the HTTP Secure Server on the WAAS Express Device" section on page 10-32.)

Step 6 In the Password field, enter the same password that you defined on the WAAS Express device.

Step 7 Click **Submit**.

Note Changing the WAAS Express credentials on the Central Manager does not change the configuration on the WAAS Express device. It affects only the credentials that are stored on the Central Manager.

To configure an external TACACS+ user on the WAAS Express device, use the following configuration commands on the WAAS Express device:

```
waas-express#config t
Enter configuration commands, one per line. End with CNTL/Z.
waas-express(config)#aaa new-model
waas-express(config)#aaa authentication login default group tacacs+
waas-express(config)#aaa authorization exec default group tacacs+
waas-express(config)#tacacs-server host host-ip
waas-express(config)#tacacs-server key keyword
```

To configure an external RADIUS user on the WAAS Express device, use the following configuration commands on the WAAS Express device:

```
waas-express#config t
Enter configuration commands, one per line. End with CNTL/Z.
waas-express(config)#aaa new-model
waas-express(config)#aaa authentication login default group radius
waas-express(config)#aaa authorization exec default group radius
waas-express(config)#radius-server host host-ip
waas-express(config)#radius-server key keyword
```

The external authentication server for TACACS+ or RADIUS must be Cisco ACS 4.x or 5.x.

Importing the Central Manager Certificate

The next step is to import the certificate from the Central Manager into the WAAS Express device.

To import the certificate, follow these steps:

Step 1 Log in to the Central Manager CLI.

Step 2 Display the administrative certificate by using the show crypto EXEC command:

```
waas-cm#show crypto certificate-detail admin

...
-----BEGIN CERTIFICATE-----
TIICezCCAeSgAwIBAgIEVwMK8zANBgkqhkiG9w0BAQUFADCBgTELMAkGA1UEBhMC
VVMxEzARBgNVBAgTCkNhbGlmb3JuaWExETAPBgNVBAcTCFNhbiBKb3NlMQ0wCwYD
VQQLEwRDTkJVMRswGQYDVQQKExJDaXNjbyBTeXN0ZW1zLCBJbmMxHjAcBgNVBAMT
FWRvYy13YWFzLWNtLmNpc2NvLmNvbTAeFw0wODA3MjQxOTMwMjNaFw0xMzA3MjMx
OTMwMjNaMIGBMQswCQYDVQQGEwJVUzETMBEGA1UECBMKQ2FsaWZvcm5pYTERMA8G
A1UEBxMIU2FuIEpvc2UxDTALBgNVBAsTBENOQlUxGzAZBgNVBAoTEkNpc2NvIFN5
c3RlbXMsIEluYzEeMBwGA1UEAxMVZG9jLXdhYXMtY20uY2lzY28uY29tMIGfMA0G
CSqGSIb3DQEBAQUAA4GNADCBiQKBgQCy10xBfsUDTh5imYwkkterx/IqkNQO7KB/
M0wqIK2j4zj4BpR1ztKaFyEtGjqGpxPBQ54V9EHGmGUljx/Um9PORk3AXyWoUsDf
o0T2Z94FL5UoVUGzUia6/xiUrPCLNf6BLBDGPQg970QtZSU+DYUqjYHzDgv6yXFt
viHARbhZdQIDAQABMA0GCSqGSIb3DQEBBQUAA4GBADKF7aIeQ+Uh4Y2zZJwlaIF7
ON+RqDvtyy4DNerEN9iLI4EFO/QJ+uhChZZU8AKR8u3OnLPSNtNck33OWwMemcOd
QGhnsMtiUq2VuSh+A3Udm+sMLFguCw5RmJvqKTrj3ngAsmDBW3uaK0wkPGp+y3+0
2hUYMf+mCrCOwBEPfs/M
-----END CERTIFICATE-----
```

Step 3 Copy the certificate text, which is the part in between the BEGIN CERTIFICATE and END CERTIFICATE lines in the output.

Step 4 Log in to the WAAS Express device CLI.

Step 5 Configure a certificate for the Central Manager:

```
waas-express#config t
Enter configuration commands, one per line. End with CNTL/Z.
waas-express(config)#crypto pki trustpoint wcm
```

```
waas-express(ca-trustpoint)#enrollment terminal pem
waas-express(ca-trustpoint)#exit
waas-express(config)#crypto pki authenticate wcm

Enter the base 64 encoded CA certificate.
End with a blank line or the word "quit" on a line by itself
```

Step 6 Paste in the certificate that you copied from the Central Manager in Step 3.

Configuring a WAAS Express Device Certificate

The WAAS Express device needs a certificate that is requested by the Central Manager when establishing HTTPS communication. This procedure describes how to configure a persistent self-signed certificate on the router, but you can also use a CA signed certificate.

To configure a WAAS Express device certificate, follow these steps:

Step 1 Log in to the WAAS Express device CLI.

Step 2 Create a self-signed certificate on the router:

Note Due to CSCsy03412, you must configure **ip domain name** *name* before enrolling the certificate. If you do not configure **ip domain name**, IOS regenerates the self-signed certificate upon reload and this affects the communication with the WAAS Central Manager.

```
waas-express#config t
Enter configuration commands, one per line. End with CNTL/Z.
waas-express(config)#crypto pki trustpoint local
waas-express(ca-trustpoint)#enrollment selfsigned
waas-express(ca-trustpoint)#subject-alt-name routerFQDN
waas-express(ca-trustpoint)#exit
waas-express(config)#crypto pki enroll local
% Include the router serial number in the subject name? [yes/no]: yes
% Include an IP address in the subject name? [no]: yes
Enter Interface name or IP Address[]: 10.10.10.25
Generate Self Signed Router Certificate? [yes/no]: yes

Router Self Signed Certificate successfully created
```

If the WAAS Express device certificate changes after the WAAS Express device is registered with the Central Manager, you must reimport the certificate into the Central Manager. For details, see the "Reimporting WAAS Express Certificate" section on page 10-34.

Enabling the HTTP Secure Server on the WAAS Express Device

The Central Manager and a WAAS Express device communicate using the HTTPS protocol. You must enable the HTTP secure server on the WAAS Express device.

To enable the HTTP secure server, follow these steps:

Step 1 On the WAAS Express device, enable the HTTP secure server:

```
waas-express#config t
Enter configuration commands, one per line. End with CNTL/Z.
waas-express(config)#ip http secure-server
```

Step 2 Configure authentication for the HTTP server for a local user as follows:

```
waas-express(config)#ip http authentication local
```

If you are using external TACACS+ or RADIUS user authentication, configure authentication for the HTTP server as follows:

```
waas-express(config)#ip http authentication aaa
```

Note If you do not configure local or AAA authentication for the HTTP server, only the enable password is used for authentication. (The default is **ip http authentication enable**, which uses only the enable password and no username.) If this default configuration is used, it is not necessary to define a username credential for the WAAS Express device on the Central Manager. (See the "Configuring a User" section on page 10-30.)

Installing a License on the WAAS Express Device

The WAAS Express device requires a license to operate the WAAS Express software.

To install a permanent WAAS license, follow these steps:

Step 1 Obtain and copy the WAAS license to a location accessible to the **license** command on the WAAS Express device.

Step 2 On the WAAS Express device, install the WAAS license:

```
waas-express#license install ftp://infra/licenses/FHH122500AZ_20100811190225615.lic
```

This example uses FTP to get and install the license but there are various options available for this command. Choose one that best suits your deployment.

Step 3 Save the running configuration:

```
waas-express#write memory
Building configuration...
[OK]
```

Configuring an NTP Server

It is important to keep the time synchronized between devices in your WAAS network. You should already have an NTP server configured for the Central Manager (see the "Configuring NTP Settings" section on page 10-5).

To configure an NTP server for the WAAS Express device, on the WAAS Express device use the **ntp server** global configuration command, as follows:

```
waas-express#config t
Enter configuration commands, one per line. End with CNTL/Z.
waas-express(config)#ntp server 10.10.10.55
```

Registering the WAAS Express Device

The final step in setting up a WAAS Express device with the Central Manager is to register the device. You will need to know the IP address of the Central Manager.

To register a WAAS Express device with the Central Manager, follow these steps:

Step 1 On the WAAS Express device, register with the Central Manager:

```
waas-express#waas cm-register https://CM_IP_Address:8443/wcm/register
```

In the URL for this command, specify the Central Manager IP address as indicated. Be sure to include a colon and the port number of 8443.

If a permanent WAAS license is not installed on the WAAS Express device, you must accept the terms of the evaluation license to continue. The evaluation license is valid for 60 days.

Step 2 Save the running configuration:

```
waas-express#write memory
Building configuration...
[OK]
```

After the successful registration of the WAAS Express device in the Central Manager, the Central Manager initially shows the device on the Manage Devices page with a management status of Pending and a license status of Active. After the Central Manager retrieves the device configuration and status, the management status changes to Online and the license status changes to Permanent (or Evaluation, Expires in x weeks y days).

Reimporting WAAS Express Certificate

If the WAAS Express device certificate changes after you have registered the WAAS Express device with the Central Manager, you must reimport a matching certificate into the Central Manager.

To reimport a WAAS Express device certificate, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Admin > Security > WAAS Express Certificate**. The Modifying WAAS Express device certificate window appears, as shown in Figure 10-2.

The Certificate Info tab shows the certificate information for the WAAS Express device. The Certificate in PEM Encoded Format tab shows the certificate in PEM format. You can copy the certificate from this tab to use in the paste operation in the next step.

Figure 10-2 Modifying WAAS Express Device Certificate Window

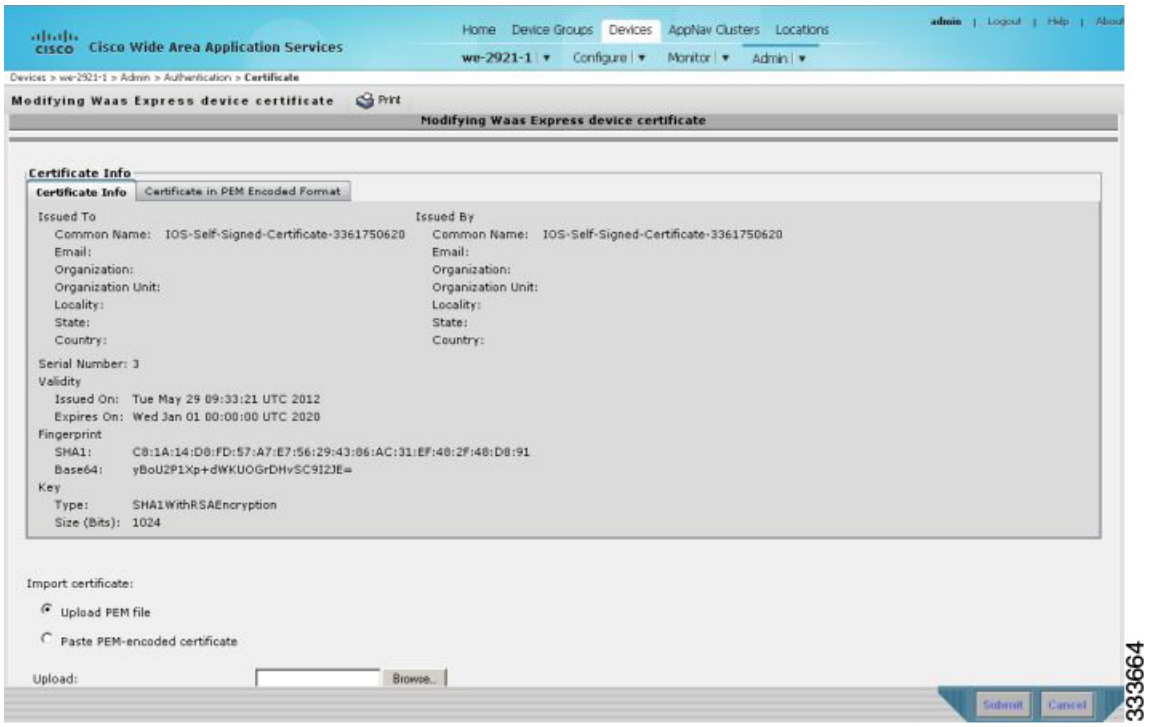

Step 3 Import this certificate into the Central Manager by selecting one of the following radio buttons that are shown in both tabs:

- **Upload PEM file**—Click Browse and locate the PEM file containing the certificate.
- **Paste PEM-encoded certificate**—Paste the PEM encoded certificate in the text field that appears.

Step 4 Click **Submit**.

CHAPTER 11

Using the WAE Device Manager GUI

This chapter describes how to use the WAE Device Manager GUI, which is a separate interface from the WAAS Central Manager GUI. The WAE Device Manager is a web-based management interface that allows you to control and monitor an individual WAE device in your network.The WAAS Central Manager device does not have a WAE Device Manager interface. In many cases, the same device settings are found in both the WAE Device Manager and the WAAS Central Manager GUI. For this reason, we recommend that you always configure device settings from the WAAS Central Manager GUI if possible.

When you change device settings in the WAE Device Manager, the changes are propagated to the WAAS Central Manager and override the group settings for that device. If you later decide that you want the group settings to override the settings that you configured from the WAE Device Manager, you can use the group override features in the WAAS Central Manager GUI. For more information, see the "Overriding Group Configuration Settings" section on page 3-7.

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

This chapter contains the following sections:

Launching the WAE Device Manager

Each WAAS device is managed separately using the WAE Device Manager web-based interface. You can launch the WAE Device Manager remotely from any location on the WAAS network using Internet Explorer (required).

To launch the WAE Device Manager, go to https://*Device_IP_Address*:8443/mgr

The Login window of the WAE Device Manager appears. Enter your username and password in the fields provided and click **Login**. The default username is admin and the default password is default.

The WAE Device Manager interface appears. (See Figure 11-1.)

Figure 11-1 WAE Device Manager Interface

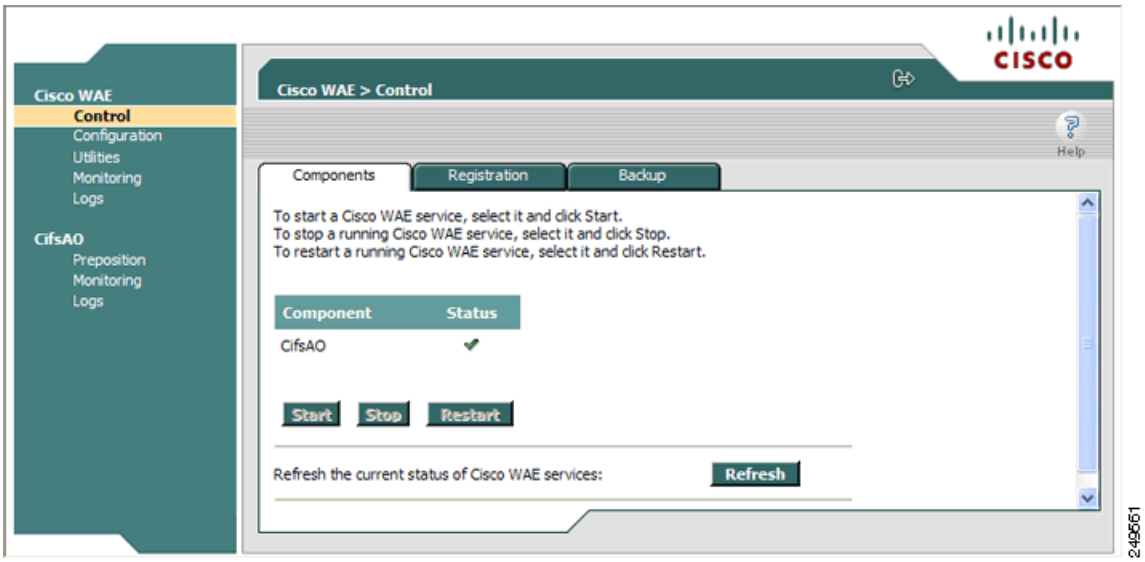

A Quick Tour of the WAE Device Manager

The WAE Device Manager is divided into two sections. The area on the left displays the navigation area. The area on the right displays information about the options that you have selected from the navigation area.

The navigation area allows you to navigate the management screens for different WAE components. The navigation area includes the following options:

- Cisco WAE—Allows you to start and stop the WAE components, register and unregister the WAE, back up and restore configuration files, and use various WAE utilities. For more information, see the "Managing a Cisco WAE" section on page 11-3.
- CifsAO—Allows you to monitor preposition tasks, view CIFS device statistics, and view the log. For more information, see the "Managing a CIFS Accelerator Device" section on page 11-19.

  The CifsAO option only appears if you have enabled the transparent CIFS accelerator on this WAAS device. For more information, see the "Enabling and Disabling the Global Optimization Features" section on page 13-3.

The options in the navigation area include suboptions, which when selected, display additional tabs in the display area. Mandatory fields in the display area are indicated with an asterisk. If you click **Save** without entering a value in a mandatory field, an error message is displayed. Click the **Back** link to return to the window where the error occurred.

Information displayed in tables can be sorted by clicking the column headers. Clicking the header a second time sorts the information in reverse order.

As you navigate in the WAE Device Manager, your current location is always displayed across the top of the display area.

To log out of the WAE Device Manager, click the  icon on the upper-right side of the display area.

Note JavaScripts, cookies, and popup windows must be enabled in the web browser to use the WAE Device Manager.

WAE Management Workflow

After WAEs have been deployed and registered (as described in the *Cisco Wide Area Application Services Quick Configuration Guide*), use the WAE Device Manager to perform the following actions:

- Start and stop components as described in the "Starting and Stopping Components" section on page 11-5.
- Register and unregister the WAE as described in the "Registering and Unregistering a WAE" section on page 11-6.
- Back up and restore configuration files as described in the "Backing Up the Configuration Files" section on page 11-6.
- Configure Windows authentication as described in the "Configuring Windows Authentication" section on page 11-10.
- Define component-specific notification recipients as described in the "Defining Notification Settings" section on page 11-15.
- Run WAE maintenance utilities as described in the "Utilities Option" section on page 11-17.
- View the details, current status, and history of preposition tasks performed on CIFS devices as described in the "Preposition Option" section on page 11-20.
- View SNMP-generated information and graphs about each WAE component as described in the "Monitoring the WAE" section on page 11-22.
- View the logs for each WAE component as described in the "Viewing WAE Logs" section on page 11-27.

Managing a Cisco WAE

You use the Cisco WAE menu item in the navigation area to perform basic operations such as viewing the status of WAE components and stop or start components on the WAE. Figure 11-2 shows the Cisco WAE Control window.

Figure 11-2 Cisco WAE Control Window

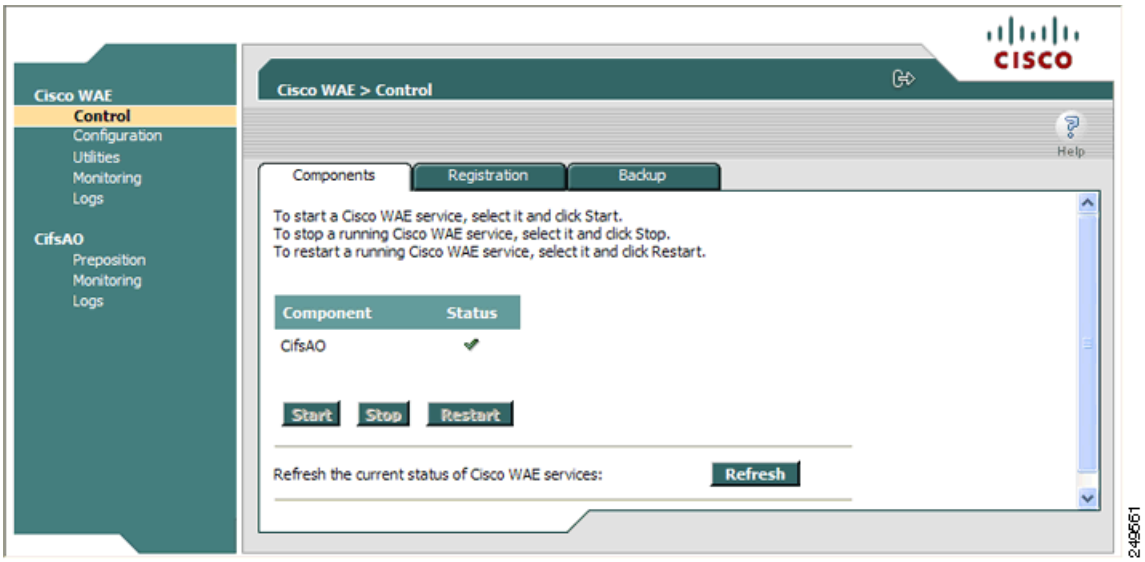

The Cisco WAE menu item includes the following options:

- Control—Enables you to control the WAE and its components as described in the "Control Option" section on page 11-4.
- Configuration—Enables you to perform basic configuration tasks as described in the "Configuration Option" section on page 11-8.
- Utilities—Enables you to run various maintenance utilities on the WAE as described in the "Utilities Option" section on page 11-17.
- Monitoring—Enables you to view tables and graphs about the CPU and disk utilization in the WAE as described in the "Monitoring the WAE" section on page 11-22.
- Logs—Enables you to view event logs for various WAE subsystems as described in the "Viewing WAE Logs" section on page 11-27.

Control Option

The Control option displays the following tabs:

- Components—Enables you to view the working status of each WAE component. You can start, stop, and restart any component. For more information, see the "Starting and Stopping Components" section on page 11-5.
- Registration—Enables you to register or unregister the WAE with the WAAS Central Manager. For more information, see the "Registering and Unregistering a WAE" section on page 11-6.
- Backup—Enables you to download and save WAE configuration files and to restore these files back to the WAE, if required. For more information, see the "Backing Up the Configuration Files" section on page 11-6 and the "Restoring the Configuration Files" section on page 11-7.

Starting and Stopping Components

The Components tab enables you to view which components are running and which components are not, and allows you to start, stop, and restart components.

From this tab you can click **Refresh** to update the status of each component and update the WAE Device Manager interface to reflect recent changes made to the device from the WAAS Central Manager GUI. For example, if the device is configured to be a transparent CIFS accelerator device while you are logged into the WAE Device Manager, that change is not reflected until you either click **Refresh** or log in again to the WAE Device Manager.

Note If a component is not running, most of its configuration can be performed offline. However, any configuration changes made to the component will take effect only after it is restarted.

Note Do not stop or start a component if the device is not registered to a WAAS Central Manager.

To start and stop components, follow these steps:

Step 1 In the Components tab of the Cisco WAE Control window, choose the component that you want to activate and click **Start**.

After a few seconds, a green checkmark appears next to the selected component, indicating its status is running, as shown in Figure 11-3.

Figure 11-3 Components Tab—Starting Components

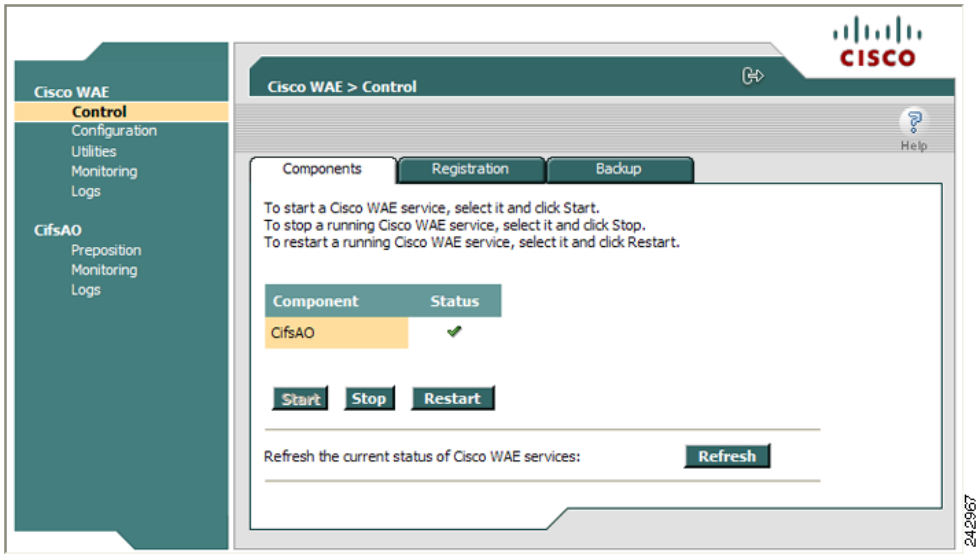

- To stop a component, choose the component from the list and click **Stop**.

  After a few seconds, a red appears next to the selected component, indicating that it is no longer running.

- To restart a WAE component, choose the component from the list and click **Restart**.

- To display the current status of the WAE components, click **Refresh**.

Registering and Unregistering a WAE

The Registration tab enables you to register the WAE with the specified WAAS Central Manager or unregister the WAE. After the WAE is registered, you can view and manage it from the WAAS Central Manager GUI.

To register the WAE, follow these steps:

Step 1 In the Cisco WAE Control window, click the **Registration** tab. (See Figure 11-4.)

Figure 11-4 Cisco WAE Control—Registration Tab

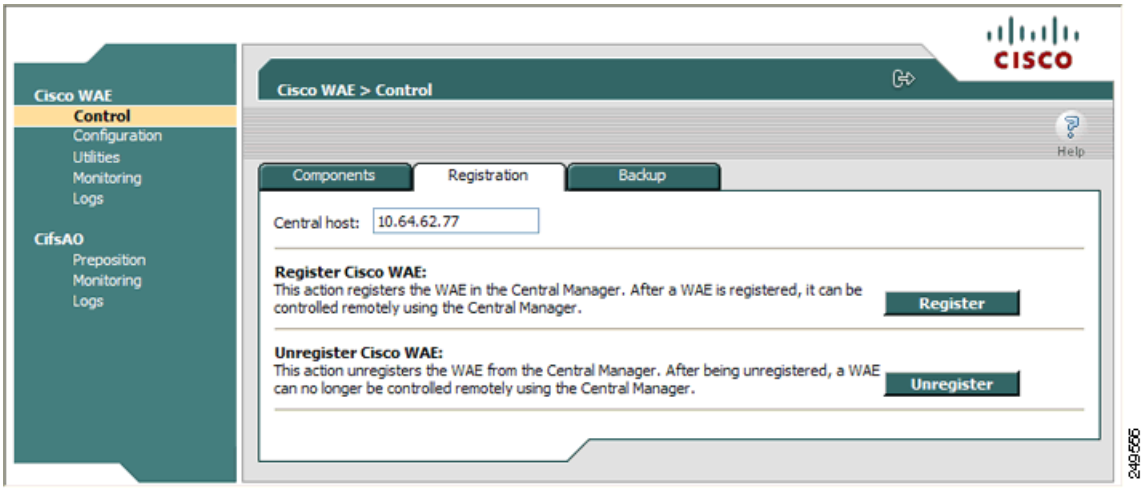

Step 2 In the Central Host field, verify that the address of the WAAS Central Manager is displayed. If no address appears in this field, then the WAE is not registered with a Central Manager.

Step 3 Click **Register** to register the WAE.

The "Registration will update the WAE properties in the WAAS Central Manager. Are you sure?" message is displayed. Click **OK**. If successful, the "Appliance registered successfully" message is displayed.

Step 4 Click **Unregister** to unregister the Cisco WAE.

If successful, the "Appliance unregistered successfully" message is displayed.

Note When you unregister a WAE, any policies defined for it in the WAAS Central Manager GUI are removed.

Backing Up the Configuration Files

The Backup tab enables you to back up and restore the configuration files of the WAE.

To back up the WAE configuration, follow these steps:

Step 1 In the Cisco WAE Control window, click the **Backup** tab. (See Figure 11-5.)

Figure 11-5 Cisco WAE Control —Backup Tab

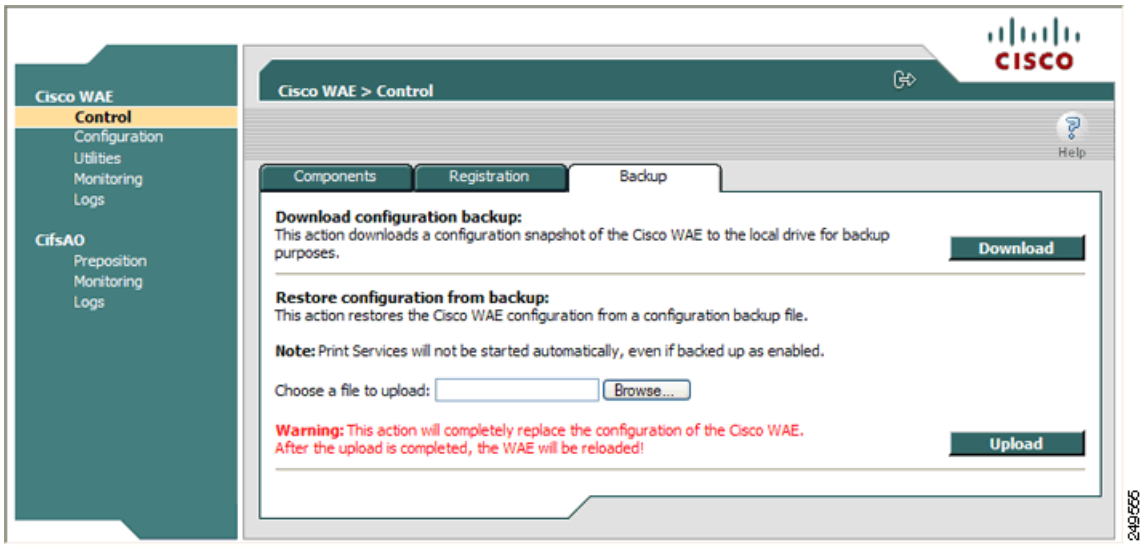

Step 2 In the Download configuration backup area, click **Download**.

Step 3 In the File Download window, click **Save**.

Step 4 In the Save As window, browse to where you want to save the file. You can also change the filename.

Step 5 Click **Save**.

The WAE configuration files are downloaded to the selected destination folder and stored in a single, compressed file.

For information about restoring files from a backup, see the "Restoring the Configuration Files" section on page 11-7.

Restoring the Configuration Files

The Backup tab enables you to restore the configuration files of the WAE. Restoring the configuration returns the WAE to its previous state when the backup was performed.

To restore the configuration files, follow these steps:

Step 1 In the Restore configuration from backup area, click **Browse** to navigate to the location of the backup file that you want to restore.

Step 2 Click **Upload** to restore the selected configuration files.

Note After the upload is completed, the WAE will be reloaded.

Configuration Option

The Configuration option for the Cisco WAE menu item displays the following tabs:

- SNMP—Allows you to enable event MIB and logging traps on the WAE. For more information, see the "Configuring SNMP Settings" section on page 11-8.
- Networking—Allows you to view WAE settings defined during initial device setup described in the *Cisco Wide Area Application Services Quick Configuration Guide*. For more information, see the "Viewing Network Settings" section on page 11-9.
- Windows Authentication—Allows you to define the settings required by the WAE for Windows authentication to enable device login and CLI configuration. For more information, see the "Configuring Windows Authentication" section on page 11-10.
- Notifier—Allows you to define the e-mail address to which notifications are sent when alerts are generated by the WAE. For more information, see the "Defining Notification Settings" section on page 11-15.

Configuring SNMP Settings

The SNMP tab allows you to configure the SNMP settings on the Cisco WAE. To configure the SNMP settings, click the **SNMP** tab in the Configuration window. The SNMP tab appears. (See Figure 11-6.)

Figure 11-6 WAE Configuration—SNMP Tab

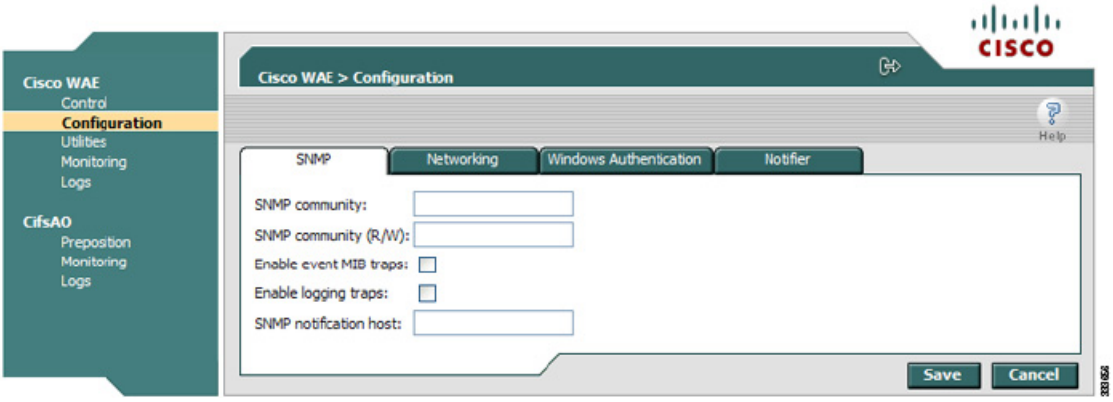

This tab allows you to configure the following settings:

- SNMP community—Sets the SNMP community string for read access, which is used as a password for authentication when accessing the SNMP agent of the WAE.
- SNMP community (R/W)—Sets the SNMP community string for read/write access, which is used as a password for authentication when accessing the SNMP agent of the WAE.

- Enable event MIB traps—Allows the WAE to send event MIB traps to the SNMP host specified in the SNMP notification host field.
- Enable logging traps—Enables logging traps on the device.
- SNMP notification host—Enter the IP address or hostname of your SNMP host so that the WAE can send MIB and logging traps to the host.

Click **Save** after making any changes to this page, or click **Cancel** to disregard your changes so that they do not take effect.

Viewing Network Settings

The Networking tab (see Figure 11-7) enables you to view the connection parameters between the WAE and the LAN.

To view the WAE connection settings, click the **Networking** tab in the Configuration window.

Figure 11-7 Cisco WAE Configuration—Networking Tab

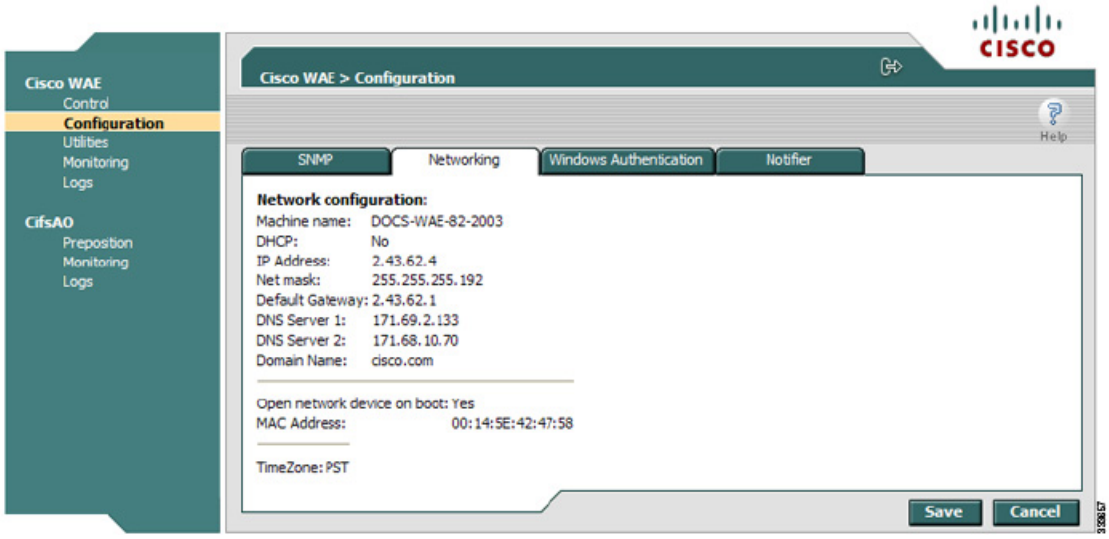

The Networking tab contains the following information:

- Network connection flags—The network status flags.
- Mode—The duplex and speed of the connection.
- Machine name—The hostname of the WAE.
- DHCP—Whether a DHCP server is available on the network.
- IP Address
- Net mask
- Default Gateway
- DNS Server 1
- DNS Server 2
- Domain Name
- MAC Address

- Time Zone

Configuring Windows Authentication

The WAAS Central Manager GUI and the WAE Device Manager use Pluggable Authentication Modules (PAM) for user login authentication. Administrative users defined in the WAAS Central Manager GUI are distributed to the WAE Device Managers. Administrative user authentication is performed only upon login to the WAAS Central Manager GUI or the WAE Device Manager. Each WAE has a default GUI and CLI user with the username admin and password default. This user account cannot be deleted, but the password can be changed.

Note In situations where the CLI user account information conflicts with the management GUI configuration, the management GUI configuration will overwrite any conflicting CLI user account information at the time of configuration distribution. A warning is displayed to CLI users after configuring CLI user account settings to inform users of this behavior.

This section contains the following topics:

Understanding Login Authentication and Authorization Through the Local Database

Local user authentication and authorization use locally configured usernames and passwords to authenticate administrative user login attempts. The login and passwords are local to each WAE.

By default, local user login authentication is enabled as the primary authentication method. You can disable local user login authentication only after enabling one or more of the other administrative login authentication methods. However, when local user login authentication is disabled, and you disable all other administrative login authentication methods, local user login authentication is reenabled automatically.

Windows Domain authentication is another user login authentication method. You can use the console, Telnet, FTP, SSH, or HTTP (WAAS Central Manager and WAE Device Manager interfaces) to authenticate Windows Domain users.

Supported Authentication Methods

When you enable Windows authentication on your WAE, you can configure additional settings that make the authentication process of your users, WAE, and services more secure when they register with the domain controller.

CIFS supports the following Windows authentication methods on the WAE:

- NTLMv2 authentication—A Windows authentication protocol that is built into most Windows operating systems.
- Kerberos—A Windows authentication protocol that uses secret-key cryptography and is built into Windows 2003 Server.

Note Windows domain authentication is not performed unless a Windows domain server is configured on the WAAS device. If the device is not successfully registered, authentication and authorization do not occur. WAAS supports authentication by a Windows domain controller running only on Windows Server 2000 or Windows Server 2003.

If you are using NTLM authentication, the Windows domain server must be installed with the option to support pre-Windows 2000 operating systems. (On the installation Permissions screen of the Windows server dcpromo wizard, select "Permissions compatible with pre-Windows 2000 server operating systems.")

LDAP Server Signing

Lightweight Directory Access Protocol (LDAP) server signing is a configuration option of the Microsoft Windows Server's Network security settings. This option controls the signing requirements for LDAP clients such as the WAE. LDAP signing is used to verify that an intermediate party did not tamper with the LDAP packets on the network and to guarantee that the packaged data comes from a known source.

The WAAS software supports login authentication with Windows 2003 domains when the LDAP server signing requirements option for the Domain Security Policy has been set to "Require signing." LDAP server signing allows the WAE to join the domain and authenticate users securely.

Note When you configure your Windows domain controller to require an LDAP signature, you must also configure LDAP server signing on the WAE from the CLI by using the **smb-conf section "global" name "ldap ssl" value "yes"** global configuration command. You cannot enable this option using the WAE Device Manager interface. For information on using the **smb-conf** command, see the *Cisco Wide Area Application Services Command Reference*.

Setting Up Windows Authentication

The Windows Authentication tab allows you to configure the security settings on the WAE.

To configure Windows Authentication, follow these steps:

Step 1 Log into the WAE Device Manager.

Step 2 In the Configuration window, click the **Windows Authentication** tab.

The Window Authentication window appears. (See Figure 11-8.)

Figure 11-8 Cisco WAE Configuration—Windows Authentication Tab

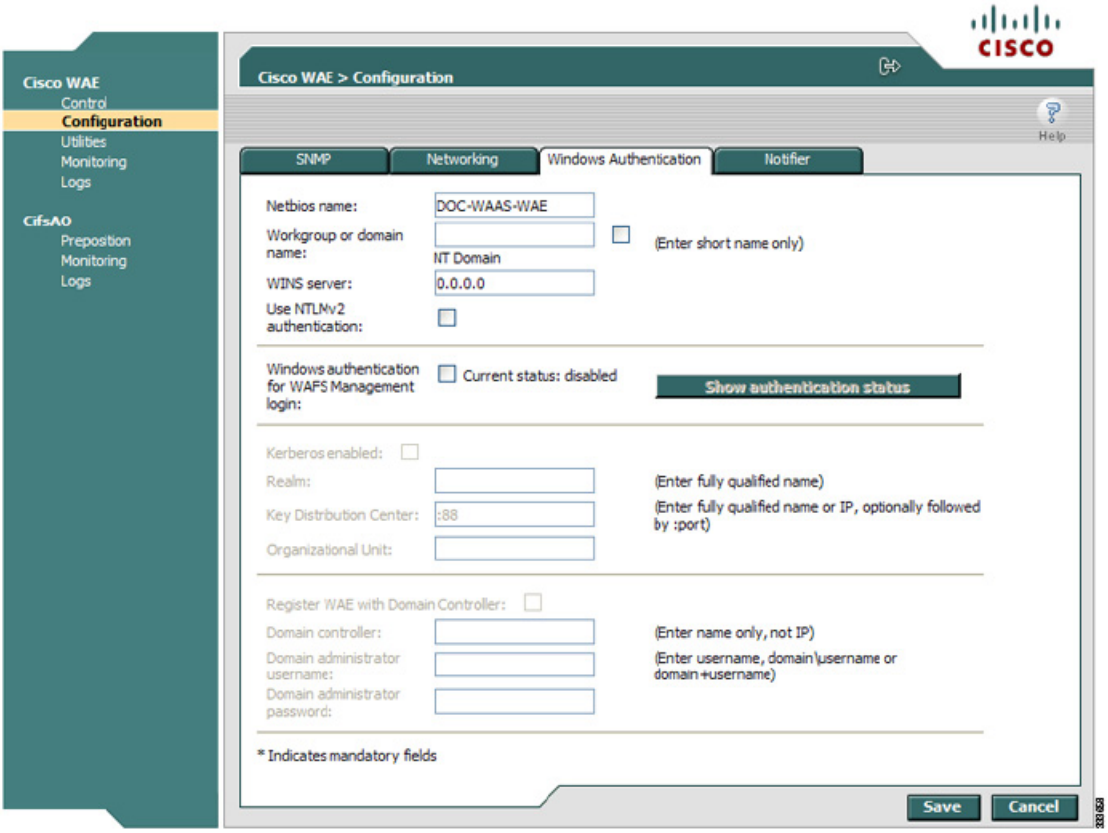

Step 3 Enter the NetBIOS name.

The NetBIOS name cannot exceed 15 characters nor contain special characters.

Note By default, the NetBIOS name field is automatically populated with the hostname of the file engine. If this hostname changes, the NetBIOS field is not automatically updated with the new name.

Step 4 Enter the workgroup or domain name in the short name format, and check the **NT Domain** check box if the workgroup/domain is a Windows NT 4 domain.

For example, if your domain name is cisco.com, the short name format is cisco. If your workgroup or domain is a Windows 2000 or Windows 2003 domain, do *not* check the **NT Domain** check box.

If the NT Domain check box is checked, the domain name and short name format can contain a period (.), but be careful not to enter the fully qualified name for the NT domain.

Step 5 Enter the IP address or hostname for the WINS server that you are using.

Step 6 Check the **Use NTLMv2 authentication** check box to enable NTLMv2 authentication.

Note Enable NTLMv2 support *only* if all clients have their security policy set to "Send NTLMv2 responses only/Refuse LM and NTLM." Using NTLM v2 when the clients do not require it could cause authentication to fail.

Step 7 Check the **Windows authentication for WAFS Management login** check box to use Windows Domain to authenticate Telnet, FTP, console, SSH, and user interface (WAAS Central Manager GUI and WAE Device Manager) logins to CIFS (WAFS).

When you add users through the WAAS Central Manager GUI, you are given the option to configure users as local users who have their login password stored on the WAE. Local users are authenticated by the WAE, but nonlocal users are commonly verified using Windows Domain authentication.

Step 8 If you are using Kerberos authentication, check the **Kerberos enabled** check box and then specify the following information:

- The fully qualified name of the Kerberos realm. All Windows 2000 domains are also Kerberos realms, but the realm name is always the all uppercase version of the domain name.
- The fully qualified name or IP address of the Key Distribution Center. You can also specify a port using the following format: *ip address* or *name*:*port number*. For example, 10.10.10.2:88.
- The organizational unit.

You can only enable Kerberos authentication if at least one of the boxes described in Step 7 is checked. After you enable Kerberos, make sure that the clock on your WAE is within 5 minutes of the clock on your domain controller. Otherwise, your domain controller will refuse to use Kerberos for authentication.

If you are using a Windows 2000 (with SP4) or Windows 2003 (with SP1) domain controller, you should enable Kerberos authentication.

Step 9 If your domain controller has been configured to require LDAP server signing, you need to use the WAAS CLI to enable LDAP server signing on the WAE by using the **smb-conf section "global" name "ldap ssl" value "yes"** global configuration command. For information on using the **smb-conf** command, see the *Cisco Wide Area Application Services Command Reference*.

Step 10 Check the **Register WAE with Domain Controller** check box.

Note You need to register the WAE with the domain controller whenever you enable or disable Kerberos, enable Windows authentication, or change the NetBios name, workgroup, or Kerberos realm.

A series of fields display under the check box. Enter the following information in these fields:

- Domain controller (enter the name, not the IP address).

  You can only enter the NetBios name of the domain controller when Kerberos is disabled. If Kerberos is enabled, you can enter the fully qualified domain name of the domain controller.
- Domain administrator username (enter the username, domain\username, or domain+username).
- Domain administrator password.

Step 11 Click **Save**.

The Windows Authentication settings are saved, and the WAE is registered with the domain controller.

Step 12 Verify if Windows Authentication is working correctly. See the "Checking the Status of Windows Authentication" section on page 11-13.

Checking the Status of Windows Authentication

After you enable Windows Authentication, you can check the status of Windows Authentication and view the results of built-in tests that can help you resolve authentication issues.

A Windows Authentication problem can occur if you incorrectly configure the settings described in the "Setting Up Windows Authentication" section on page 11-11. Problems can also occur if the configuration of your domain controller changes.

The Authentication Details window shows the following information:

- A list of winbind Authentication tests
- The results of each test
- A pass or fail indicator
- Troubleshooting tips to help you resolve why a test failed

To check the status of Windows Authentication, follow these steps:

Step 1 On the Windows Authentication tab, click **Show authentication status**.

A message appears that explains the authentication status could take a while to display and that the WAE's performance could be impacted while the authentication status is being obtained.

Step 2 In the message dialog box, click **OK** to proceed or click **Cancel** to not display the authentication details.

If you clicked OK, the Authentication Details window appears. (See Figure 11-9.)

Figure 11-9 Authentication Details Window

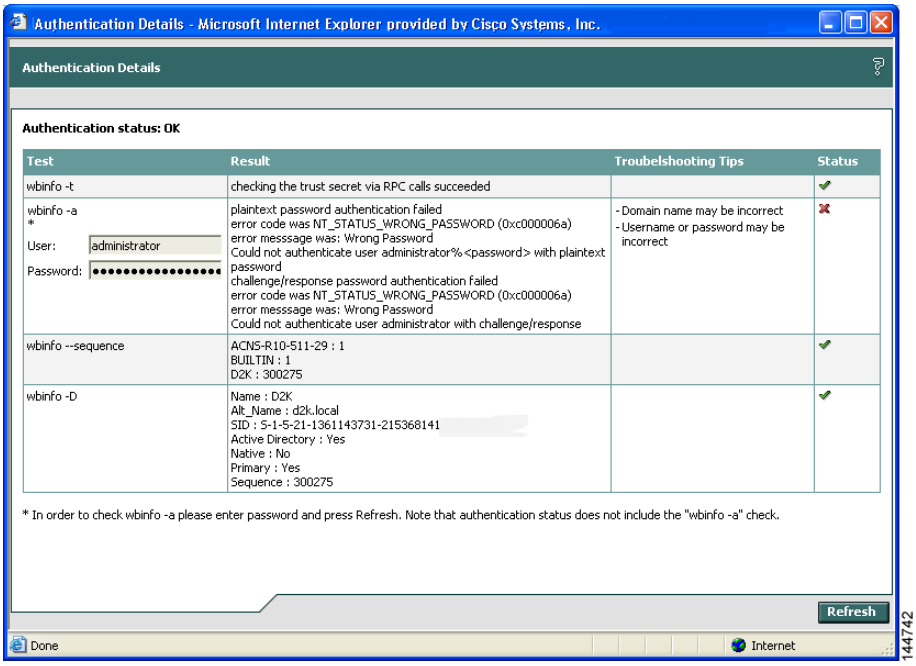

Step 3 Check the Authentication status field at the top of the window.

If the status field displays "OK," then Windows Authentication is functioning correctly. If this field displays "Not OK," then proceed to the next step.

Step 4 View the status of each test, and resolve any failures using the provided troubleshooting tips.

Table 11-1 describes these tests.

Table 11-1 Authentication Test Descriptions

Test	Description
wbinfo -t	Verifies that the workstation trust account created when the Samba server is added to the Windows domain is working.
wbinfo -a	Tests the domain credentials based on the specified username and password. To run this test, enter the appropriate username and password, and then click **Refresh**. Wait for the test results to be displayed.
wbinfo -D	Shows information from Samba about the domain.
wbinfo --sequence	Shows the sequence numbers of all known domains.
Time skew	Shows the time offset between the WAE and the KDC server. The time offset must be within 5 minutes; otherwise, the Windows KDC server refuses to use Kerberos for authentication. You can use the WAAS CLI to configure the time on the WAE. This test is performed only when Kerberos authentication is enabled.

Step 5 Click **Refresh** to ensure that all the tests complete successfully.

Defining Notification Settings

The Notifier tab allows you to define the e-mail address to which notifications are sent when alerts are generated by the WAE.

To define notification settings, follow these steps:

Step 1 In the Configuration window, click the **Notifier** tab. (See Figure 11-10.)

Figure 11-10 Notifier Tab

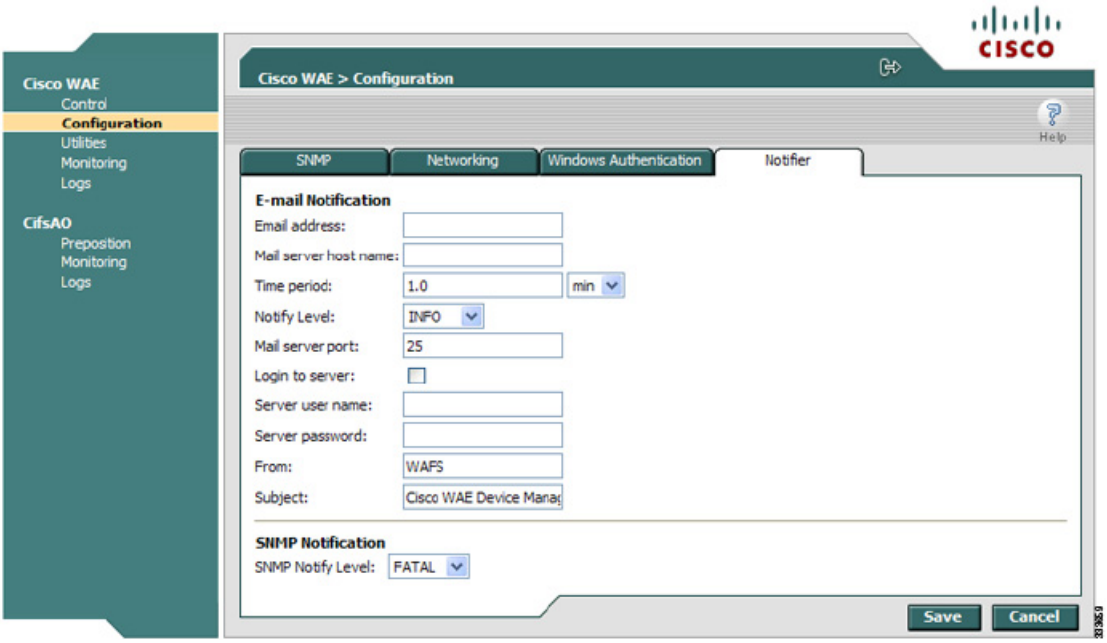

Step 2 In the Email address field, enter the address to which notifications about this WAE are sent.

Step 3 In the Mail server host name field, enter the name of the mail server host.

Step 4 In the Time period field, enter the time interval for notifications to accumulate until they are sent through e-mail and choose the relevant time unit from the drop-down list (min or sec).

Step 5 From the Notify Level drop-down list, choose the minimum event severity level for generating notifications.

Step 6 In the Mail server port field, enter the port number for connecting with the mail server.

Step 7 Check the **Login to server** check box if the WAE must log in to the mail server to send notifications. If this option is selected, additional fields are enabled.

Step 8 In the Server username field, enter the username for accessing the mail server.

Step 9 In the Server password field, enter the password for accessing the mail server.

Step 10 In the From field, enter the text that should appear in the From field of each e-mail notification.

Step 11 In the Subject field, enter the text that should appear as the subject of each notification.

Step 12 From the SNMP Notify Level drop-down list, choose the minimum event severity level for generating SNMP notifications.

Step 13 Click **Save**.

Utilities Option

The Utilities option displays the following tabs:

- Support—Allows you to dump WAE data to an external location for support purposes. For more information, see the "Running Support Utilities" section on page 11-17.
- WAFS Cache Cleanup—Allows you to remove all files from the CIFS (WAFS) cache. For more information, see the "Running the Cache Cleanup Utility" section on page 11-18.
- File Server Rename—Allows you to rename a file server in the CIFS (WAFS) cache. For more information, see the "Running the File Server Rename Utility" section on page 11-19.

Running Support Utilities

The Support tab displays product information about the WAE, including the WAAS software version and build number running on the device.

The Support tab also allows you to download a system report that provides a snapshot of the current state of the WAE and its operation, including the configuration log files of various components. You can send this report to Cisco Technical Support (TAC) if you need assistance.

Note Downloading a full system report can impact the performance of the WAE. For this reason, we recommend downloading the system report during nonpeak hours or limiting the date range of the report.

To download the system report, follow these steps:

Step 1 In the Utilities window, click the **Support** tab.

The Support window appears. (See Figure 11-11.)

Figure 11-11 Utilities—Support Tab

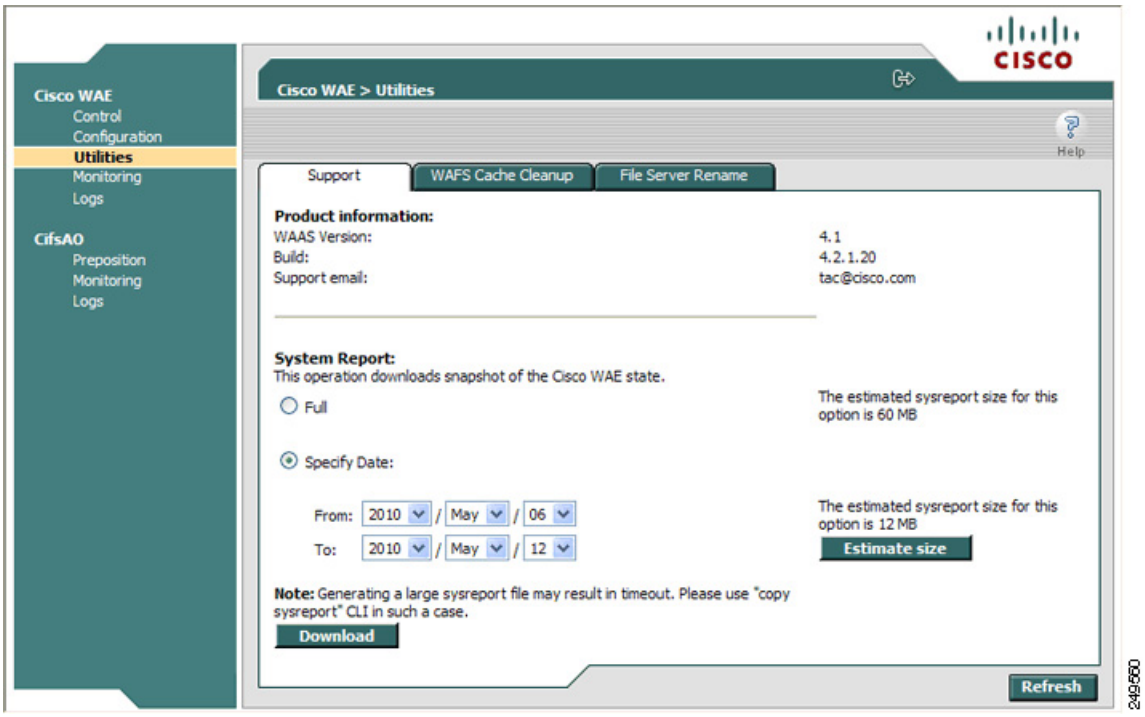

Step 2 In the System Report area, choose one of the following radio buttons:

- **Full** to download a full system report.
- **Specify Date:** to download a report for the time range that you specify (default is the past 7 days).

Step 3 Click **Estimate size** to view the size of the report.

The actual size of the report may vary from the estimate. If the estimated size is large, you may want to specify a smaller time frame and download multiple smaller reports to minimize the stress on the WAE.

Step 4 Click **Download**.

A message informs you that downloading the report can affect the performance of all services on the device.

Step 5 Click **OK** to start the collection process.

Step 6 In the File Download window, click **Save**.

Step 7 In the Save As window, browse to where you want to save the file. (You can also change the filename.) Click **Save**. The file is saved in tar gzip format.

Running the Cache Cleanup Utility

The WAFS Cache Cleanup tab enables you to remove all files from the CIFS device cache.

To run the cache cleanup utility, follow these steps:

Step 1 In the Utilities window, click the **WAFS Cache Cleanup** tab.

The WAFS Cache Cleanup window appears. (See Figure 11-12.)

Figure 11-12 Utilities—WAFS Cache Cleanup Tab

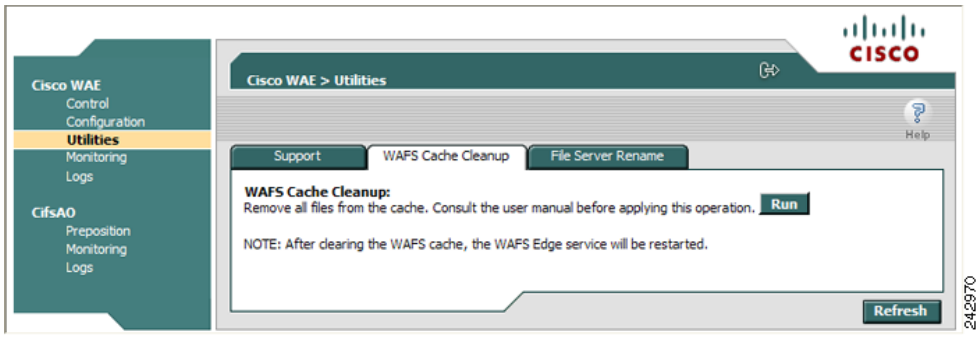

Step 2 Click **Run** to erase the contents of the cache.

Running the File Server Rename Utility

The File Server Rename tab enables you to change the resource location for all resources of a given file server name on the WAAS device. This function changes the file server name for the files in the CIFS cache.

To run the file server rename utility, follow these steps:

Step 1 If the CifsAO component is running, stop it as described in the "Starting and Stopping Components" section on page 11-5.

Step 2 In the Utilities window, click the **File Server Rename** tab.

Step 3 In the Current File Server name field, enter the current name.

Step 4 In the New File Server name field, enter the new name and click **Run** for the new name to take effect.

Note Do not specify the name of another existing cached file server in the New File Server name field. If you do specify an existing name as the new name, the cached contents of this file server are overwritten with the cached contents of the file server you are renaming.

Managing a CIFS Accelerator Device

The CifsAO option in the navigation area allows you to monitor preposition tasks, view CIFS device statistics, and view the log. The CifsAO option appears only if you are using transparent CIFS accelerator mode.

The CifsAO option includes the following menu items:

- Preposition—Allows you to monitor the progress of preposition policies created in the WAAS Central Manager GUI. In addition, you can optionally terminate preposition tasks. For more information, see the "Preposition Option" section on page 11-20.

- Monitoring—Allows you to view CIFS (WAFS) device statistics in tables and graphs as described in the "Monitoring the Cisco WAE Component" section on page 11-24.
- Logs—Allows you to view the event log related to the CIFS accelerator. For more information, see the "Viewing Cisco WAE Logs" section on page 11-28.

Preposition Option

The Preposition option allows you to view the details and current status of preposition policies created in the WAAS Central Manager GUI. These policies define which files are proactively placed in the WAAS device cache according to a prearranged schedule. Prepositioning enables system administrators to strategically place large, frequently accessed files at the network edge during off-peak hours, increasing efficiency and providing end users with quick first-time access of those files.

You can view information such as the root directory containing the files being prepositioned, the schedule for each policy, and the status of the most recent task for each policy. You can also view a detailed task history for each policy, and manually terminate any tasks in progress.

To view preposition policies for this device, follow these steps:

Step 1 In the navigation area, click **Preposition**.

The CifsAO > Preposition window appears. (See Figure 11-13.)

Figure 11-13 CifsAO Preposition Window

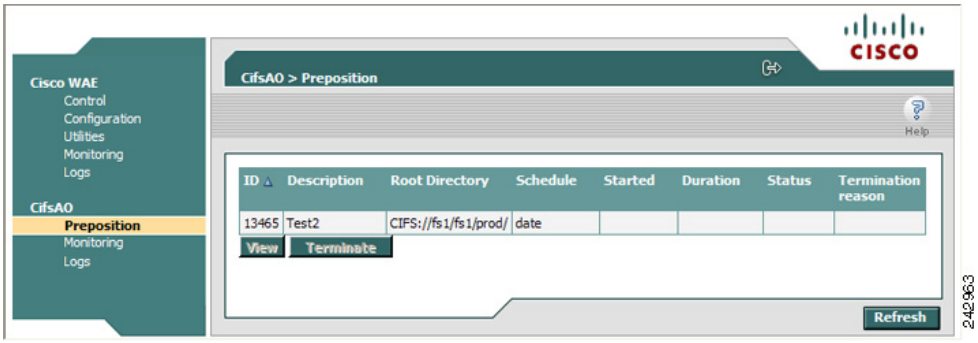

The Preposition window contains a table that displays all the preposition policies assigned to this CIFS Edge device. For each policy, the following information is displayed:

- ID—ID number of the selected policy.
- Description—Descriptive name assigned to the policy.
- Root Directory—Source directory for the content being prepositioned.
- Schedule—Defined schedule for the policy.
- Started—When this policy was last invoked by the system.
- Duration—Elapsed time of the latest task.
- Status—Current status of the policy, updated every time the refresh button is clicked. If the task defined by the policy is currently being run, its status is In Progress. A preposition task in progress can be terminated.
- Termination reason—Reason the policy was terminated.

Step 2 Choose a policy in the table and click **View** to view a detailed task history (iterations of a selected policy).

The Preposition Task Details window appears. (See Figure 11-14.)

Figure 11-14 Preposition Task Details Window

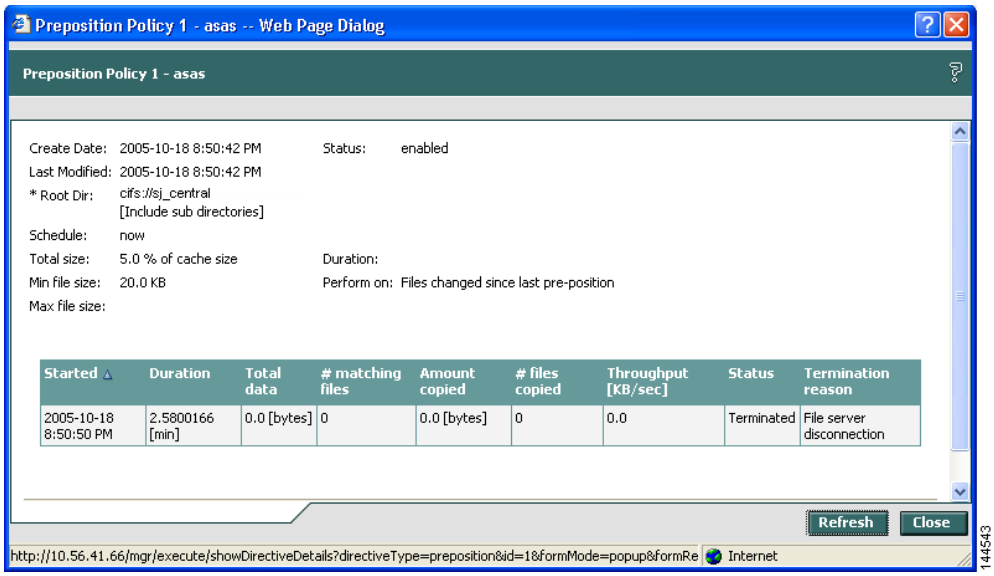

The top half of the Preposition Policy window displays the following details about the selected policy:

- Create Date—When the policy was created.
- Last Modified—When the policy was last modified.
- Total size—Limit placed on the total size of the files being prepositioned, if any.
- Min file size—Minimum size of files in the root directory (and subdirectories if they are part of the preposition policy) that are affected by the policy.
- Max file size—Maximum size of files in the root directory (and subdirectories if they are part of the preposition policy) that are affected by the policy.
- Perform on—Which files to preposition from the selected location—those files that have changed since the last preposition, those files changed during a defined interval, or all files.

The lower half of the Preposition Policy window contains a table that displays the most recent tasks performed by the selected policy (up to the last 10 iterations), including the following information:

- Total data—Total amount of data to be transferred by the policy.
- # matching files—Number of files matching the defined filter of the policy.
- Amount copied—Total amount of data copied by the policy during its most recent run. (This amount may be less than the amount in the Total data field if the policy is currently in progress, or if the policy did not complete its run, for example, due to time constraints placed on its operation.)
- # files copied—Number of files copied by the policy during its most recent run.
- Throughput—Throughput achieved by the policy in kilobits per second (Kbps).
- Termination reason—Reason that the policy was terminated, if relevant. Policies can be terminated due to time or space constraints placed on the policy or to a decision by the administrator to manually terminate its operation.

Step 3 Click **Close** to return to the Policies window.

Note To update the information displayed in the Policies window, click **Refresh**.

Terminating a Preposition Task

You can terminate a preposition task that is in progress at any time. This action does not delete the preposition policy that generated the task; the system will still perform the task described by the policy when the next scheduled time arrives.

Note Do not terminate a preposition task if the device is not registered to a WAAS Central Manager.

To terminate a preposition task, follow these steps:

Step 1 In the Policies window, select a preposition policy with a status of In Progress and click **Terminate**. A confirmation message is displayed.

Step 2 Click **Yes** to terminate the task. If you click View to display the Preposition Policy window, the table that displays the task history contains a message indicating that the latest task was terminated by the administrator.

Monitoring the WAE

The Monitoring option available for the Cisco WAE and transparent CIFS accelerator components enables you to view detailed tables that describe the current state of the WAE. It also provides graphs that display historical data about the selected components. These graphs enable you to track WAE statistics for a day, week, month, or an entire year.

Note WAE statistics and graphs are generated by the freeware MRTG utility. For details, go to http://people.ee.ethz.ch/~oetiker/webtools/mrtg/.

The monitoring options differ for each WAE component as described in Table 11-2.

Table 11-2 Monitoring Options by Component

Component	Monitored Statistics
Cisco WAE	CPU and disk drive utilization
Transparent CIFS accelerator	CIFS traffic and cache

This section contains the following topics:

- Monitoring a Transparent CIFS Accelerator, page 11-25

Monitoring Graphs

The WAAS software generates four historical graphs for each monitored statistic. Each graph presents a different range of time for the selected data as follows:

- Daily—Displays data for the past 24 hours. Each data point represents a 5-minute average.
- Weekly—Displays data for the past seven days. Each data point represents a 30-minute average.
- Monthly—Displays data for the past five weeks. Each data point represents a 2-hour average.
- Yearly—Displays data for the past 12 months. Each data point represents a one-day average.

The maximum value over the given time period and the current value for the statistic being monitored is also displayed below each of these graphs.

Viewing Options

You can view an index window of the daily graphs for all the monitored statistics available for a component, or you can view the four historical graphs for a particular statistic (for example, cache utilization) at once.

Figure 11-15 shows a sample screen when a user chooses to view the index graphs.

Figure 11-15 Sample Index Graph Window

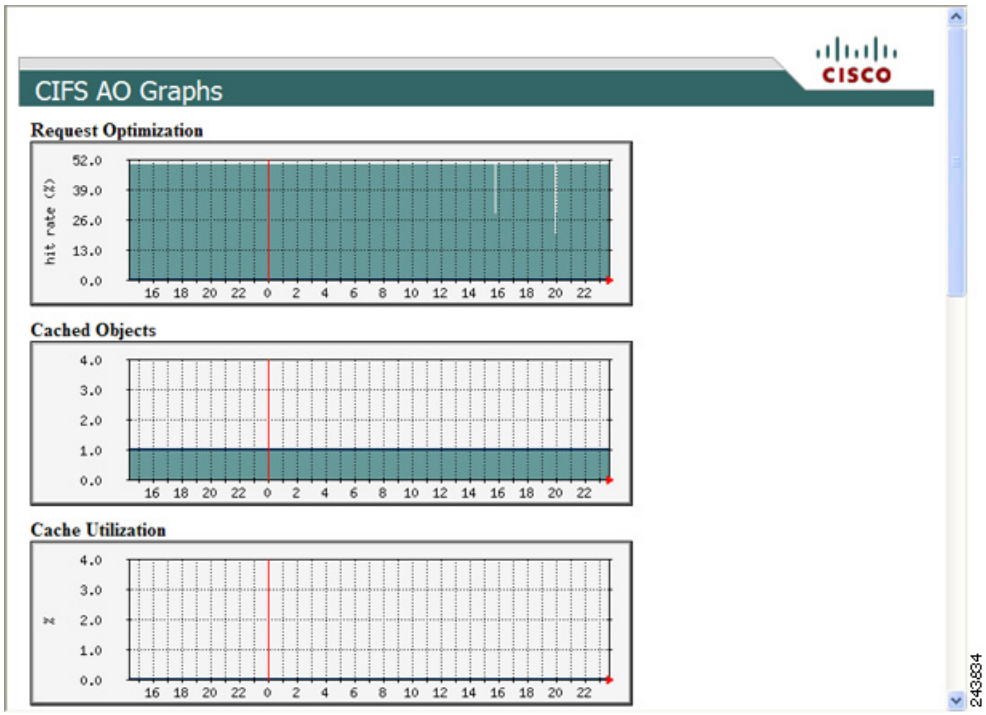

Tip Each graph in an index window acts as a link. Clicking on the graph displays all four historical graphs for the selected statistic. For example, clicking the Request Optimization graph in the index graphs window displays the daily, weekly, monthly and yearly Request Optimization historical graphs. Clicking the Back button in the browser returns you to the index graphs.

Figure 11-16 shows a sample screen when a user chooses to view the historical graphs for a particular statistic.

Figure 11-16 Sample Historical Graph Window

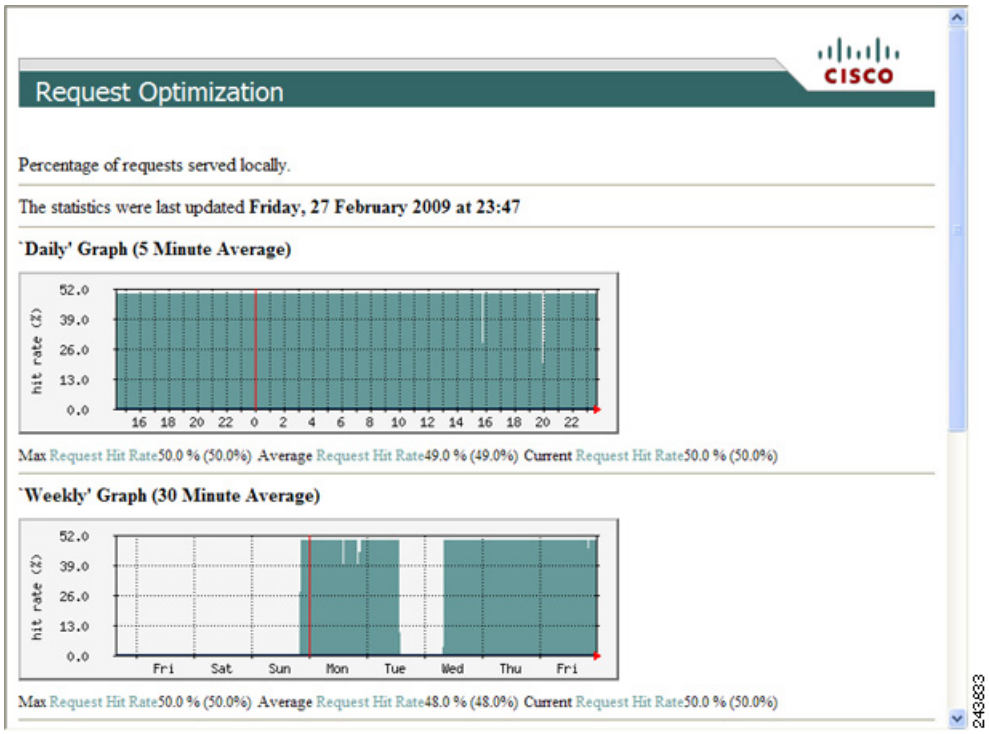

Note Graphs can be printed using the Print command in your browser.

Monitoring the Cisco WAE Component

The Monitoring option for the Cisco WAE component displays a table with the statistics monitored on a WAE. From this table, you can display historical graphs that indicate the central processing unit (CPU) utilization and disk drive utilization on the WAE.

CPU utilization is a measure of the amount of bandwidth used by the CPU versus the total bandwidth available. The amount is expressed as a percentage. Disk drive utilization is a measure of the amount of disk space that is being used on all disk drives versus the total disk space available. This amount is also expressed as a percentage.

To monitor the WAE component, follow these steps:

Step 1 In the navigation area, click **Monitoring** under the **Cisco WAE** menu item.

The Cisco WAE Monitoring window appears. (See Figure 11-17.)

Figure 11-17 Cisco WAE Monitoring Window

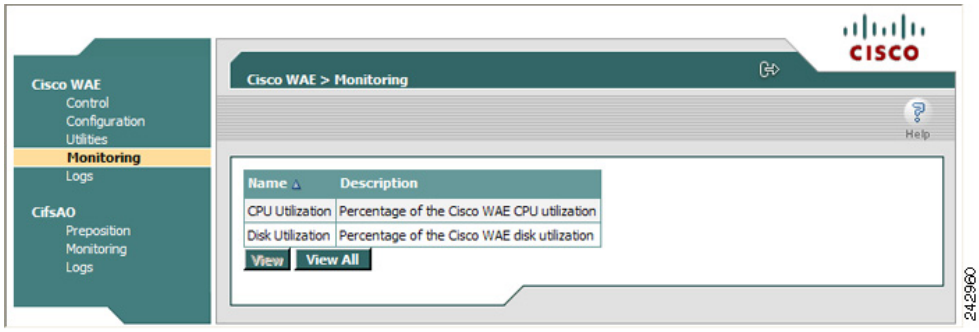

Step 2 Do one of the following:

- Choose the statistic that you want to view (by clicking in its row), and then click **View** to display a popup window that contains the historical graphs for that statistic.
- Click **View All** to display the index window with the daily graphs for both statistics on the WAE component.

Monitoring a Transparent CIFS Accelerator

The Monitoring option displays the following tabs:

- CIFS—Displays data about the status of the CIFS protocol and the selected device.
- Cache—Displays data about the device cache.
- Graphs—Displays a list of graphs that are available for the device.

Note The SNMP parameters displayed in the CIFS and Cache tabs are contained in a special MIB file.

To monitor a transparent CIFS accelerator follow these steps:

Step 1 In the navigation area, click **Monitoring** under the **CifsAO** menu.

The Monitoring window appears and the CIFS tab is displayed.

The CIFS tab displays the following CIFS-related information:

- Total Time Saved—Total time saved by CIFS acceleration.
- Total KBytes read—Total number of kilobytes read by clients (both through the cache and remotely) from this device using the CIFS protocol.
- Total KBytes written—Total number of kilobytes written by clients to this device using the CIFS protocol.

- Remote requests count—Total number of client CIFS requests that were forwarded remotely over the WAN. The name of this statistic is a link that you can use to display its historical graphs (without first going to the Graphs tab). Local requests are also shown on these graphs.
- Local requests count—Total number of client CIFS requests handled locally by this device. The name of this statistic is a link that you can use to display its historical graphs (without first going to the Graphs tab). Remote requests are also shown on these graphs.
- Total remote time—Total amount of time, in milliseconds, spent by this device to process all client CIFS requests that were sent remotely over the WAN.
- Total local time—Total amount of time, in milliseconds, spent by this device to process all client CIFS requests that were handled locally.
- Connected sessions count—Total number of CIFS sessions connected on this device. The name of this statistic is a link that you can use to display its daily, weekly, monthly, and yearly graphs (without first going to the Graphs tab).
- Open files count—Total number of open CIFS files on this device. The name of this statistic is a link that you can use to display its daily, weekly, monthly, and yearly graphs (without first going to the Graphs tab).
- CIFS Command Statistics—Table of statistics on CIFS commands. For each command type, the table lists the total number of requests, the number of remote requests, the number of asynchronous requests, the average time in milliseconds spent by this device to process each request that was handled locally, and the average time in milliseconds spent by this device to process each request that was sent remotely over the WAN.

  To reset the CIFS statistics, click the **Reset CIFS Statistics** button below the table.

Step 2 Click the **Cache** tab.

The Cache tab displays the following information:

- Maximum cache disk size—Maximum amount of disk space (in gigabytes) allocated to the CIFS device cache.
- Current cache disk usage—Current amount of disk space (in kilobytes) used by the CIFS device cache. The name of this statistic is a link that you can use to display its historical graphs (without first going to the Graphs tab).
- Maximum cache resources—Maximum number of resources (files and directories) allowed in the CIFS device cache.
- Current cache resources—Current number of resources contained in the CIFS device cache. The name of this statistic is a link that you can use to display its historical graphs (without first going to the Graphs tab).
- Evicted resources count—Number of resources that have been evicted from the cache since the device was started.
- Last eviction time—Time when a cache eviction last occurred.
- Cache size high watermark—Percentage of disk usage that causes the CIFS device to begin evicting resources.
- Cache size low watermark—Percentage of disk usage that causes the CIFS device to stop evicting resources.
- Cache resources high watermark—Percentage of total cache resources that causes the CIFS device to begin evicting resources.
- Cache resources low watermark—Percentage of total cache resources that causes the CIFS device to stop evicting resources.

- Last evicted resource age—Amount of time that the last-evicted resource spent in the CIFS device cache.
- Last evicted resource access time—Last time that the last-evicted resource was accessed.

Viewing WAE Logs

You can view event information logged by the Cisco WAE and the CifsAO components. The event information available varies based on the component that you are viewing.

This section contains the following topics:

- WAE Logs, page 11-27
- Viewing Cisco WAE Logs, page 11-28

WAE Logs

You can configure what you want displayed for each log file and save the log to a file locally as described in the following sections:

- Setting Display Criteria, page 11-27
- Viewing Log Entries, page 11-28
- Saving Log File Information, page 11-28

Setting Display Criteria

All WAE logs allow you to set the criteria for the data that you want to display as shown in Figure 11-18.

Figure 11-18 WAE Log Data Criteria

To set the criteria for viewing log information, follow these steps:

Step 1 Choose the beginning date (year, month, and day) and time (hour and minutes using a 24-hour clock format) from the From drop-down list.

Step 2 Choose the ending date (year, month, and day) and time (hour and minutes using a 24-hour clock format) from the **To** drop-down list.

Step 3 (Optional) Choose the minimum severity level of events to display from the Log Level drop-down list.

By choosing the minimum severity level, all events with a severity level greater than that specified are displayed. The default is **All**.

Step 4 (Optional) Choose the number of events (one per line) to appear on a single page of the log from the **Lines** drop-down list.

The default is 100 events.

Step 5 (Optional) Enter a filter string by which the log can be further filtered.

Step 6 Click **Update**.

Viewing Log Entries

Each log entry contains the date and time that the event occurred, the severity level of the event, and a description containing the log message. The log message format varies based on the type of event.

The severity level of an event indicates the seriousness of the event. Six choices are defined and provide the follow information:

- All—Displays events of all severity levels.
- Debug—Indicates events have occurred that match those specified for debugging purposes.
- Info—Indicates an event occurred regarding the proper operation of the component. No user action is required with this type of event.
- Warning—Indicates a minor problem occurred on a component. The component should be able to overcome the incident without user intervention.
- Error—Indicates a problem occurred that affected the proper operation of the component. User intervention is likely required.
- Fatal—Indicates a severe problem occurred on a component that may have caused it to stop operating. User intervention is required.

Saving Log File Information

You can save a log as a text file and download it to your local drive.

To save a log as a text file, follow these steps:

Step 1 Set up your log with the date range and time frame that you want to save, using the **From** and **To** drop-down lists. (See the "Setting Display Criteria" section on page 11-27.)

Step 2 Set up the severity level of the events you want to view.

For more information, see the "Setting Display Criteria" section on page 11-27.

Step 3 Click **Update**.

Step 4 Click **Download**.

The File Download window appears.

Step 5 Click **Save** in the File Download window.

Step 6 Specify the directory where you want to save the log file.

Step 7 Click **OK**.

Viewing Cisco WAE Logs

Each WAE component generates its own log files.

The Cisco WAE component generates these logs:

- Manager log—Displays events related to the WAE Device Manager and WAAS Central Manager GUI components, such as configuration changes and WAE registrations and notifications that other WAE components were started or stopped.
- WAFS Watchdog log—Displays events related to the watchdog utility, which monitors the other application files inside the WAE and restarts them, if necessary.

The CIFS accelerator generates one log that displays all events related to CIFS accelerator operation.

To view Cisco WAE and CIFS accelerator logs, follow these steps:

Step 1 In the navigation area, click the **Logs** option under the Cisco WAE or CifsAO component.

Figure 11-19 shows the Logs window for the Cisco WAE component.

Figure 11-19 Cisco WAE Component Logs Window

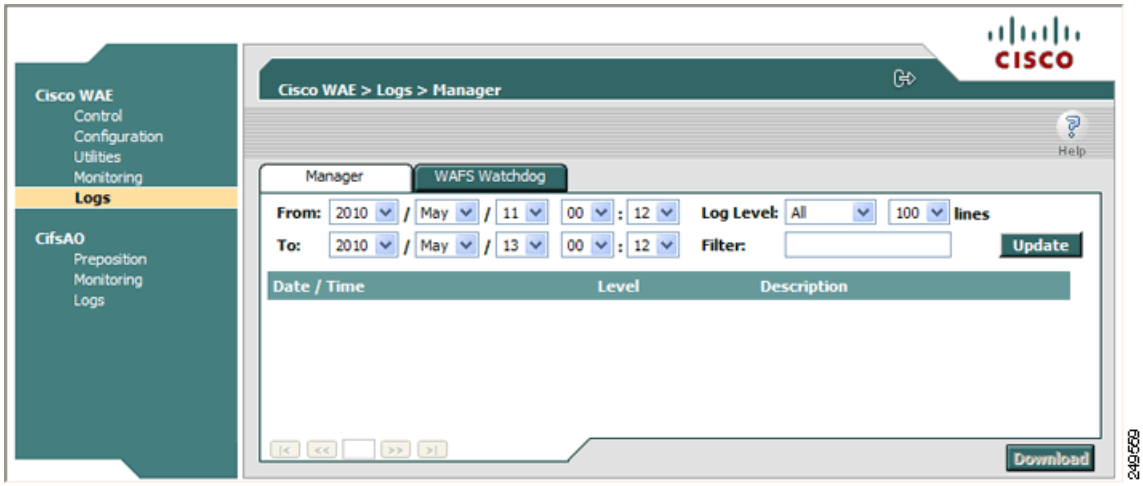

Step 2 If you selected the Cisco WAE, click the **Manager** or **WAFS Watchdog** tab to choose the log that you want to view.

Step 3 Set up your display criteria using the **From**, **To**, **Level**, and **Lines** drop-down lists. (See the "Setting Display Criteria" section on page 11-27.)

Step 4 (Optional) Set a filter on the log so that only events containing specific words or phrases are displayed by entering the relevant free text in the **Filter** text box.

Step 5 Click **Update**. The Logs window is refreshed according to your selected criteria.

Note Navigation arrows (|< << 1 >> >|) appear at the bottom of each log window when the number of events is greater than the number of lines selected per window.

CHAPTER 12

Configuring File Services

This chapter describes how to configure file services, which allows branch office users to more efficiently access data stored at centralized data centers. The file services feature overcomes the WAN latency and bandwidth limitations by caching data on Edge WAEs near branch office users. WAAS file services uses either the CIFS or SMB application accelerators.

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

This chapter contains the following sections:

About File Services

Enterprises today have remote offices in different parts of the country and around the world. Typically, these remote offices have their own file servers to store and manage the data needed by their local users.

The problem with this method of operation is that it is costly to purchase, manage, and upgrade file servers at each remote office. A great deal of resources and manpower must be dedicated to maintaining these file servers, and especially to protect the data in case of server failure. To achieve the required level of data assurance, the remote office must devote resources to back up the data at the remote site and physically move it to a secure location, often at a considerable distance from the site. If you multiply this scenario by tens, hundreds, and thousands of remote offices, and you can see that this approach to enterprise data management not only raises costs exponentially, it also greatly increases the risks to critical data.

The logical solution in this scenario is to move all of the enterprise's important data to a central location containing the facilities, trained personnel, and storage mass required to manage the data properly. By having a data center provide backup and other storage management facilities, the enterprise can achieve better utilization of both personnel and storage, as well as a higher level of data assurance and security.

The WAN between the enterprise's data center and its remote offices tends to be unreliable and slow, with limited bandwidth and high latency. In addition, the WAN creates other obstacles to the implementation of the data center solution.

One obstacle is created by the file server protocols that operate over the WAN. Common Internet File System (CIFS), which is the file server protocol for Windows, was designed to operate over a LAN. Every file operation generates several exchanges of protocol messages between the client and the file server. This situation is usually not noticeable on the LAN, but quickly causes high latency over the WAN. Occasionally, this high latency breaks the file server protocol altogether.

Even in cases where the file server protocol is managing to function correctly over the WAN, there are typically long delays between each transaction. These delays can often cause timeouts in user applications such as word processing programs, image editing programs, and design tools, which stops them from functioning correctly.

All of these problems—unreliable WANs, file system protocol compatibility, and user application compatibility—contribute to an unfriendly work environment that negatively affects the user experience and diminishes productivity.

The WAAS file services feature overcomes the WAN latency and bandwidth limitations by caching data on Edge WAEs near the user. This data caching method allows branch office users to access centralized data at LAN-like speeds over the WAN. The solution is based on several key concepts:

- Use the WAN as little as possible—By minimizing the number of operations that need to traverse the WAN, WAAS effectively shields users from many of the obstacles that WANs create.
- Use the WAN optimally—The file services feature uses sophisticated caching, compression, and network optimization technologies, which enable the system to use the WAN optimally.
- Preserve file system protocol semantics—Although WAAS software uses its own proprietary protocol over the WAN, it leaves the complete semantics of the standard file system protocol commands intact. This is essential to preserve the correctness and coherency of the data in the network.
- Make the solution transparent to users—The best solutions are the ones that do their jobs unnoticed, without interfering with end users' operations or forcing users to change their ways of doing business. The WAAS file services solution does not require any software installations, either on the server side or at the client, and does not require the user to learn anything new. Users derive all the benefits of having a secure data center without needing to change any of their work habits.

By using the WAAS file services feature, enterprises can consolidate their file servers to a data center that provides the facilities, IT personnel, and storage devices required to manage the data properly. Figure 12-1 shows a typical deployment scenario after WAAS file services have been set up.

Figure 12-1 WAAS File Services Solution

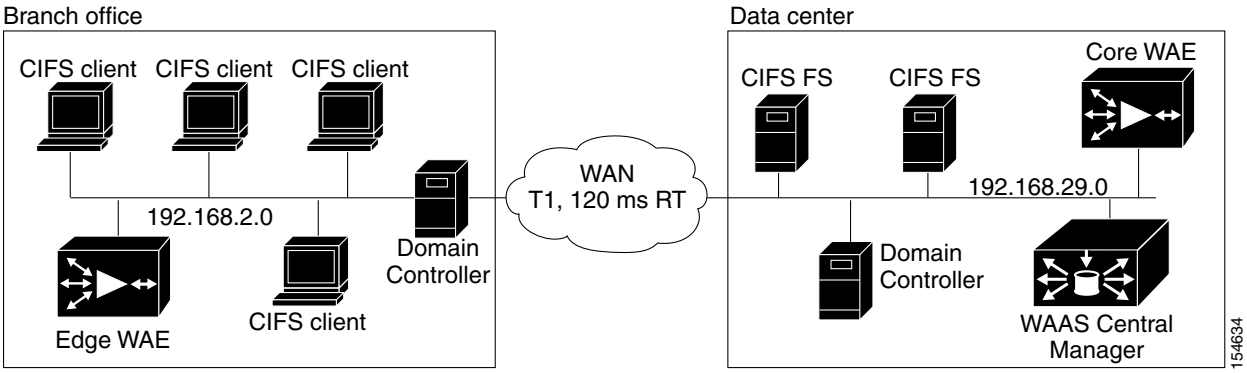

Overview of File Services Features

This section provides an overview of the WAAS file services features and contains the following topics:

To accelerate CIFS traffic, you can use one of the following two accelerators:

- CIFS—The CIFS accelerator was introduced in WAAS version 4.1.1, relies on automatic discovery, transparently accelerates CIFS traffic, supports prepositioning of files, and requires no configuration. This accelerator also supports the Windows Print accelerator, which accelerates print traffic between clients and a Windows print server. The CIFS accelerator is enabled by default.

  Supports the SMB 1.0 protocol for CIFS traffic.

- SMB—The SMB accelerator, introduced in WAAS version 5.0.1, relies on automatic discovery, transparently accelerates CIFS traffic, and does not support prepositioning or the Windows Print accelerator. This accelerator has configuration options that you can fine-tune for specific needs.

  Supports the SMB 1.0, 2.0, and 2.1 protocols for CIFS traffic and signed SMB traffic.

The CIFS and SMB accelerators are not compatible and only one can be enabled on a WAE. Enabling one automatically disables the other.

Peer WAEs must both use the same accelerator (CIFS or SMB) because the two different accelerators do not interoperate. They can coexist in the same WAAS network, but only on separate devices that are not peers.

Note Legacy mode WAFS is no longer supported beginning with WAAS version 4.4.1. Legacy WAFS users must migrate to the CIFS or SMB accelerator before upgrading.

Automatic Discovery

The automatic discovery feature allows you to enable CIFS without having to register individual file servers in the WAAS Central Manager. With the automatic discovery feature, WAAS attempts to automatically discover and connect to a new file server when a CIFS request is received.

Data Coherency

WAAS software ensures data integrity across the system by using two interrelated features – *coherency*, which manages the freshness of the data, and *concurrency*, which controls the access to the data by multiple clients.

Maintaining multiple copies of data files in multiple locations increases the likelihood that one or more of these copies will be changed, causing it to lose consistency or coherency with the others. Coherency semantics are used to provide guarantees of freshness (whether the copy is up-to-date or not) and the propagation of updates to and from the origin file server.

The WAAS software applies the following coherency semantics to its built-in coherency policies:

- Strict CIFS behavior for intra-site—Users of the same cache are always guaranteed standard, strict CIFS coherency semantics.
- Cache validation on CIFS open—In CIFS, the File Open operation is passed through to the file server. For coherency purposes, WAAS software validates the freshness of the file on every file open, and invalidates the cached file if a new version exists on the file server.

  WAAS software validates data by comparing the time stamp of a file in the cache to the time stamp of the file on the file server. If the time stamps are identical, the cached copy on the Edge WAE is considered valid and the user is permitted to open the file from the Edge WAE cache.

  If the time stamps are different, the Edge WAE removes the file from its cache and requests a fresh copy from the file server.
- Proactive cache updating—WAAS software supports the use of change notifications in CIFS environments as a way to keep cached data on the Edge WAEs up-to-date.

  When a client makes a change to a directory or file, the Edge WAE sends a change notification to the file server. The file server then sends to all the Edge WAEs a change notification that includes a list of the modified directories and files. Upon receiving the change notification, each Edge WAE checks its cache and invalidates the directories and files listed in the notification, and then updates its cache with the latest versions.

  For example, if a user edits an existing Word document and saves the changes to the Edge WAE cache, the Edge WAE sends a change notification to the file server so it knows that the file has been modified. The Edge WAE then sends the changed sections to the file server, and the file server proactively sends change notifications to the other Edge WAEs in the network. These Edge WAEs then update their cache so the file is consistent across all access points.

  This process also applies when you rename a directory, add a new subdirectory, rename a file, or create a new file in a cached directory.
- Flush on CIFS close—In CIFS, the File Close operation forces all write buffers to be flushed to the file server, and the Close request is only granted after all updates have been propagated to the file server. From a coherency standpoint, the combination of validate on file open and flush on file close ensures that well-behaved applications, such as Microsoft Office, operate in session semantics. The Open, Lock, Edit, Unlock, and Close commands are guaranteed to work correctly on the WAAS network.
- Age-based validation on directories (CIFS)—Directories are associated with a preconfigured age. When the age expires, the Edge WAE cache revalidates the directory.

  When a user first attempts to view the contents of a directory, the Edge WAE enables the file server to perform the authorization check using the directory's access control list (ACL), which contains the user and group permissions. The Edge WAE monitors which directories the user has accessed and whether the file server permitted that access. If the user tries to access the same directory again during a short period of time (aging period), the Edge WAE does not contact the file server and instead uses the cached permissions to determine if the user should be provided access. After the aging period expires, the Edge WAE contacts the file server to refresh the cached permission of the user.

  This authorization process prevents users from accessing directories and files in the cache that they do not have permission to access on the file server.

Data Concurrency

Concurrency control is important when multiple users access the same cached data to read, or write, or both. Concurrency control synchronizes this access by establishing and removing file system locks. This file-locking feature ensures data integrity and provides the following benefits:

- Enables a client to aggressively cache file data so it does not have to rely on retrieving data from the remote file server.
- Provides a performance boost in many applications running on existing CIFS client implementations.
- Preserves data integrity because only one user at a time can make changes to a section of a file.

WAAS software supports the CIFS oplocks feature, which allows a user to lock a file so the user can safely read and write data to its local cache instead of using network bandwidth to perform these functions over the WAN on the file server. By using oplocks, a user can proactively cache read-ahead data because it knows that no other user is accessing the file so there is no chance the cached data can become stale. The user can also write data to its local cache and does not need to update the file server until it closes the file or until another user requests to open the same file.

Oplocks only applies to files. The file server does not grant oplock requests on directories and named pipes.

File-Locking Process

When a user opens a file, it sends a lock request to the file server. The Edge WAE intercepts and forwards all lock requests from the user to the file server as well as all responses from the file server to the user. If no other user has a lock on the file, the file server grants an exclusive lock request so that the user can safely cache the file.

If a second user requests to open the same file, the following actions occur:

1. The file server revokes the exclusive file lock obtained by the first user.
2. The first user performs the following actions:
   - Flushes any file changes stored in its cache to the file server. This action ensures that the second user opening the file receives the latest information from the file server.
   - Deletes any of its read-ahead buffers for the file because that data is no longer guaranteed to remain up-to-date now that a second user will open the file.
3. The file server allows the second user to open the file.

Prepositioning

The prepositioning feature allows system administrators to proactively “push” frequently used files from the central storage into the cache of selected Edge WAEs. This operation provides users with faster first-time file access, and makes more efficient use of available bandwidth. You create preposition directives from the WAAS Central Manager GUI.

When an end user attempts to open a file that is not found in the Edge WAE cache, the Edge WAE retrieves it across the WAN from the file server where it is stored. Prepositioning is a feature that allows administrators to push large, frequently accessed files from file servers to selected Edge WAE caches according to a predefined schedule. Through the proper use of prepositioning, administrators can allow

users to benefit from cache-level performance even during first-time access of these files. Prepositioning improves WAN bandwidth utilization by transferring heavy content when the network is otherwise idle (for example, at night), which frees up bandwidth for other applications during the day.

The WAAS Central Manager GUI allows administrators to create multiple, overlapping preposition policies (each with its own schedule), a list of target Edge WAEs, and defined time and size constraints.

Prepositioning includes the ability to configure multiple roots. See the "Creating a New Preposition Directive" section on page 12-12.

Note Only the CIFS accelerator supports prepositioning.

Microsoft Interoperability

The WAAS file services feature interoperates with these Microsoft CIFS features:

- Active Directory for user authentication and authorization
- Offline folders in Microsoft CIFS
- Microsoft DFS infrastructure
- Windows shadow copy for shared folders, as described in the "Windows Shadow Copy for Shared Folders" section on page 12-6)

Windows Shadow Copy for Shared Folders

WAAS file services support the Shadow Copy for Shared Folders feature that is part of the Windows Server 2003/2008 operating system. This feature uses the Microsoft Volume Shadow Copy Service to create snapshots of file systems so that users can easily view previous versions of folders and files.

In a WAAS environment, users view shadow copies in the same way they would in a native Windows environment by right-clicking a folder or file from the cache and choosing **Properties > Previous Version**.

For more information about Shadow Copy for Shared Folders, including the limitations of the feature, refer to your Microsoft Windows Server 2003/2008 documentation.

Users can perform the same tasks when accessing a shadow copy folder on the Edge WAE as they can in the native environment on the file server. These tasks include the following:

- Browsing the shadow copy folder
- Copying or restoring the contents of the shadow copy folder
- Viewing and copying files in the shadow copy folder

The Shadow Copy for Shared Folders feature does not support the following tasks:

- Renaming or deleting a shadow copy directory
- Renaming, creating, or deleting files in a shadow copy directory

Supported Servers and Clients

WAAS supports Shadow Copy for Shared Folders on the following file servers:

- Windows Server 2008 and Windows Server 2008 R2
- Windows Server 2003 (with and without SP1)

- NetApp Data ONTap versions 6.5.2, 6.5.4, 7.0, and 7.3.3
- EMC Celerra versions 5.3, 5.4, and 5.6

WAAS supports Shadow Copy for Shared Folders for the following clients:

- Windows 7
- Windows Vista
- Windows XP Professional
- Windows 2000 (with SP3 or later)
- Windows 2003

Note Windows 2000 and Windows XP (without SP2) clients require the Previous Versions Client to be installed to support Shadow Copy for Shared Folders.

Preparing for File Services

Before enabling file services on your WAEs, ensure that you complete the following tasks:

- If you want to configure multiple devices with the same settings, ensure that you have created a device group that contains all the devices you want to enable with file services. For information on creating device groups, see Chapter 3, "Using Device Groups and Device Locations."
- Identify the file servers that you want to export, and refer to Table 12-1 to verify that these file servers can operate with WAAS software. Other file servers may operate with WAAS, but only those listed in the table were tested. The file server must support opportunistic locking (oplocks) and CIFS notifications.

Note The CIFS application accelerator does not support file servers that use the FAT32 file system. You can use the policy rules to exclude any FAT32 file servers from CIFS accelerator optimization.

Table 12-1 Tested File Servers

Vendor	Product	Version
Dell	PowerVault	715N
Network Appliance	FAS3140	ONTAP 7.3.3
	FAS940	ONTAP 7.0.1R.1
	FAS270	ONTAP 7.0.1R.1
	FAS250	ONTAP 7.0.1R.1
	F760	6.5.2R1P16
	F85	6.4.5
EMC	Celerra NS702	5.4.17.5
	Celerra NS702	5.4.14-3
	Celerra NS700	5.6.42-5
	Celerra NS501	5.3.12-3

Table 12-1 Tested File Servers

Vendor	Product	Version
Dell	PowerVault	715N
Microsoft	Windows NT 4.0	
	Windows Server 2000	No service pack, SP1, SP3, and SP4
	Windows Server 2003	No service pack, SP1, SP2, and R2
	Windows Server 2008[1]	SP1 and R2
Novell[2]	6.5	SP-3
RedHat	Samba	3.0.1.4a

1. With Windows 7 and Vista clients, the CIFS accelerator transparently uses the SMB1 protocol.
2. WAAS supports Novell 6.5 for CIFS optimization, server consolidation, and generic network acceleration for NCP, eDirectory/NDS, and iPrint. If your Novell file server uses the NFAP option, WAAS can optimize your Novell traffic at the transport layer as well as at the protocol layer using the WAAS CIFS adapter. NFAP is Novell's Native File Access Pack that uses the CIFS protocol on top of Novell's NCP (Novell Core Protocol).

Note Certain combinations of operating systems and file systems on a file server can result in the server responding with different timestamp precision for different SMB commands. In this situation, you may not get the highest possible CIFS optimization if the CIFS application accelerator avoids using cached files with mismatched timestamps in favor of preserving data coherency.

Using File Services on the NME-WAE

If you are running WAAS on a network module that is installed in a Cisco access router, there are specific memory requirements for supporting file services. The NME-WAE must contain at least 1 GB of RAM to support file services:

If you try to enable file services and the device does not contain enough memory, the WAAS Central Manager will display an error message.

You can check the amount of memory that a device contains in the Device Dashboard window. For details, see the "Device Dashboard Window" section on page 17-8.

Configuring File Services

To accelerate CIFS traffic, you can enable and configure either the CIFS or the SMB accelerators, as described in the following topics:

Configuring the CIFS Accelerator

The CIFS accelerator relies on automatic discovery and transparently accelerates CIFS traffic with no configuration needed.

Table 12-2 provides an overview of the steps that you must complete to configure the CIFS accelerator.

Table 12-2 Checklist for Configuring CIFS Accelerator

Task	Additional Information and Instructions
1. Prepare for file services.	Provides the tasks that you need to complete before enabling and configuring file services on your WAAS devices. For more information, see the "Preparing for File Services" section on page 12-7.
2. Enable CIFS acceleration.	Enables the transparent CIFS accelerator. For more information, see the "Enabling and Disabling the Global Optimization Features" section on page 13-3.
3. (Optional) Identify dynamic shares.	Identifies the dynamic shares on an exported file server. If your file server uses Access Based Enumeration (ABE) to give users different views of the share, you must configure the dynamic shares on the WAAS Central Manager. For more information, see the "Creating Dynamic Shares for the CIFS Accelerator" section on page 12-9.
4. (Optional) Create a preposition directive.	Defines which files are proactively copied from an exported file server to the Edge WAE cache. For more information, see the "About Preposition Directives" section on page 12-11.

Creating Dynamic Shares for the CIFS Accelerator

Many file servers use dynamic shares, which allow multiple users to access the same share but then be automatically mapped to a different directory based on the user's credentials. Dynamic shares are most commonly used on file servers to set up user home directories. For example, a directory named Home can be set up as a dynamic share on a file server so each user accessing that share is automatically redirected to their own personal directory.

If a file server contains a dynamic share or is using Access Based Enumeration (ABE), you must register that dynamic share with the WAAS Central Manager as described in this section.

Defining a dynamic share in the WAAS Central Manager allows each user to see a different view of the share and allows the operation of ABE if it is configured on the Windows Server.

Note Dynamic share configuration on the WAAS Central Manager overrides any dynamic share configuration set up directly on the WAE device using the CLI.

Before adding a dynamic share, note the following limitations:

- Each dynamic share on a file server must be unique.
- You cannot add a dynamic share if that share has a preposition directive. You must remove the preposition policy before you can add the dynamic share.
- You can use the WAAS Central Manager GUI to define any directory as a dynamic share. However, if a directory is not set up as a dynamic share on the file server, all users will read or write the same content from the same directory and will not be redirected to different directories based on their credentials.

To add a dynamic share for CIFS accelerator, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Configure > CIFS File Services > Dynamic Shares**.

A list of dynamic shares appears. The Dynamic Shares window shows all the dynamic shares configured. From this window, you can perform the following tasks:

- Edit the configuration of an existing dynamic share by clicking the **Edit** icon next to the share. You can delete the dynamic share, or modify any of the dynamic share settings.
- Add a new dynamic share definition, as described in the next steps.

Step 2 Click the **Create New Dynamic Share** icon in the taskbar to add a new dynamic share. The Creating a new Dynamic Share window appears.

Step 3 In the Name field, enter a name for the dynamic share.

The following characters are not supported in the dynamic share name: / \ : * ? " < > |
From the Assigned Domain drop-down list, choose the WAAS domain that you want to assign to the dynamic share. Only administrators who are also assigned to this WAAS domain have permission to edit the dynamic share configuration. The domain does not affect client's access to the dynamic share.

Note A WAAS domain is not the same as a DNS domain or Windows domain. For more information on WAAS domains, see the "Working with Domains" section on page 8-14.

This kind of WAAS domain does not use entities. When defining the WAAS domain, choose **None** for the Entity Type. The WAAS domain must be assigned to each WAAS admin user who needs to edit the dynamic share configuration (see the "Assigning a Domain to a User Account" section on page 8-15).

Step 4 In the File Server field, enter the name or IP address of the file server with the dynamic share.

If you specify the file server name, the edge WAE resolves it to an IP address.

The registered file servers are displayed in a drop-down list.

Step 5 In the User name, Password, and Confirm Password fields, enter the username and password credentials for the file server. If the username is in a Windows domain, specify the domain name as part of the User name field, as follows: domain\username.

These credentials are used only to access the file server when you click the **Browse** button.

Step 6 In the Share Name field, specify the location of the dynamic share by doing one of the following tasks:

- Enter the name of the dynamic share on the file server. The following characters cannot be used in the share name: \ / : * ? " < > |
- Click **Browse** next to the **Share Name** field to navigate to the correct root directory.

Note The Browse button appears only if you have at least one WAE device with the CIFS accelerator enabled and registered to the WAAS Central Manager.

Step 7 Ensure that the status of the share is set to enabled. If you change the status to disabled, the share will not be set up as a dynamic share in your WAAS environment.

Step 8 Click **Submit**.

The specified directory now functions as a dynamic share on the Edge WAE cache.

About Preposition Directives

A preposition directive allows you to determine which files should be proactively copied from CIFS file servers to the cache of selected Edge WAEs. Prepositioning enables you to take advantage of idle time on the WAN to transfer frequently accessed files to selected WAEs, where users can benefit from cache-level performance even during first-time access of these files.

Prepositioning is supported on automatically discovered file servers in the transparent CIFS accelerator

When defining a preposition directive, you select the Edge WAEs that you want to be prepositioned with content from the file server, then specify the root directories on the file server to be prepositioned. Initially, the preposition directive is in the unscheduled state. You must create a schedule that determines when and how often the content is prepositioned. Because content can be prepositioned on a regular basis, you can specify whether each new iteration of the task should copy all designated files, or only those files that have changed over a specified time interval.

In addition, you can specify time and size limits to prevent a preposition task from consuming too much bandwidth on the WAN or too much space on the Edge WAE cache. We strongly recommend that you use these limits to optimize network efficiency and prevent misuse of this feature.

When the activation time of a preposition directive arrives, a preposition task starts on the Edge WAE. Each preposition task can be monitored in the WAAS Central Manager GUI during and after processing. You can also terminate active preposition tasks if required.

Prepositioning requires that the username and password needed to access the file server be specified. These items are specified directly in the Creating New Preposition Directive window, as described in the following procedure.

Note When preposition updates are sent to the Central Manager, if any preposition file server credentials cannot be decrypted, all further preposition updates are not sent from the WAE to the Central Manager and decryption failure error messages are logged in errorlog/cms_log.current. You must reconfigure the preposition credentials from the CLI.

Prepositioning includes the ability to configure multiple roots. See the "Creating a New Preposition Directive" section on page 12-12.

When using prepositioning, both branch and data center WAEs are required (the same as for any other accelerated traffic). The branch WAE retrieves prepositioned files through an optimized connection. Verify that you have connectivity between the following network entities:

- Client to branch WAE
- Branch WAE to data center WAE
- Branch WAE to file server
- Data center WAE to file server

You will need to change any ACLs that might be blocking prepositioning traffic.

Note Though preposition directives can be created and managed by using the CLI, we recommend that you use the Central Manager GUI because you can manage prepositioning for groups of WAEs from the Central Manager. If you mix GUI and CLI configuration, unpredictable results can occur because changes on one device can affect other devices.

The following topics describe how to create and manage a preposition directive:

Creating a New Preposition Directive

To create a preposition directive, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Configure > CIFS File Services > Preposition**.

The Preposition Directives window appears. This window displays the following information about preposition directives that exist on the system:

- Preposition Directive—Name of the preposition directive.
- Type—Whether the preposition directive affects all files (Full) or just those that have changed since the last preposition task (Differential).
  - When the type is Full, all the files that match the other filters of the task and that are found on the file server are sent to the Edge to be compared with the cache.
  - When the type is Differential, only the files that are found as changed since the last successful preposition are sent to the Edge cache. The time of the last successful preposition is taken from the Edge device, so ensure that the clock is synchronized with the file server. The first scan is always a full scan. If you change the preposition task, the last successful scan time is reset.
  - When the type is Since, only the files that are found as changed within a specified time period are sent to the Edge cache.
- Status—Whether the preposition directive is enabled or disabled.
- File Server—Name of the exported file server.

From the Preposition Directive window, you can perform the following tasks:

- Edit the configuration of an existing preposition directive by clicking the **Edit** icon next to the directive. You can then delete the preposition directive, or modify any of the settings.
- Add a new preposition directive, as described in the following steps.

Step 2 Click the **Create New Preposition Directive** icon in the taskbar to create a new preposition directive.

The Creating New Preposition Directive window appears. (See Figure 12-2.)

Figure 12-2 Creating a New Preposition Directive Window

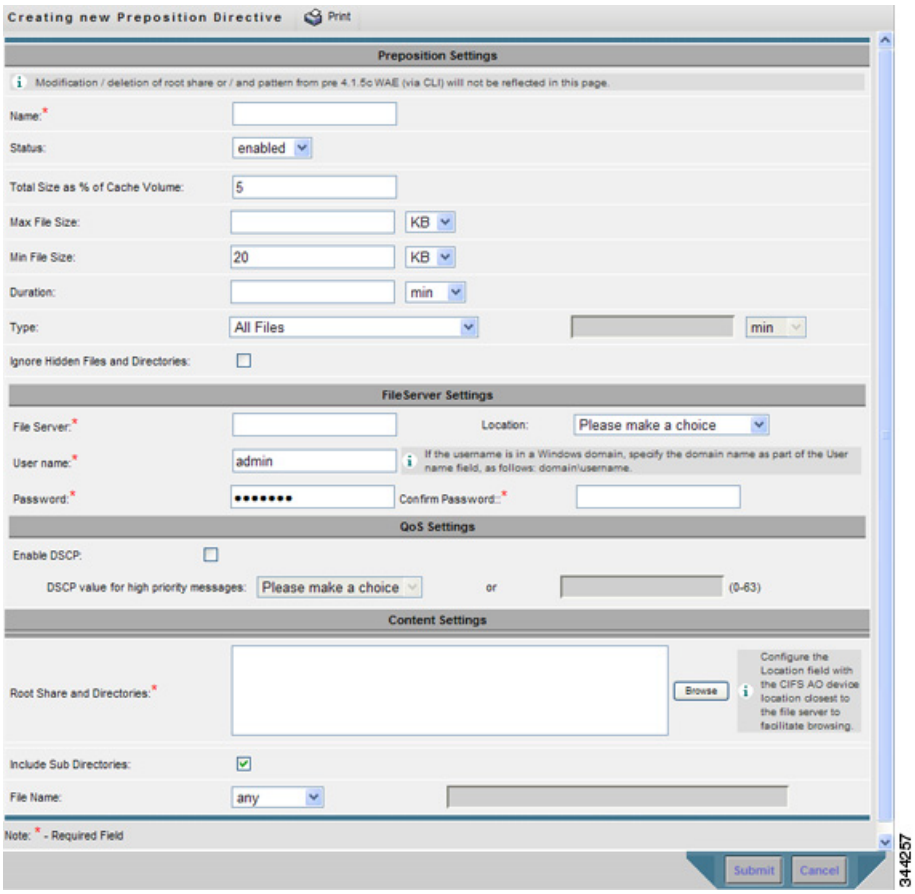

Step 3 Enter a name for the directive. The double quote (") character is not allowed in the name.

Step 4 From the Status drop-down list, choose either **enabled** or **disabled**. Disabled directives are not put into effect.

Step 5 (Optional) Define the time and size limitations using the provided fields.

Table 12-3 describes the time and size limitation fields.

Table 12-3 Preposition Time and Size Limitations

Field	Description
Total Size as % of Cache Volume	Percentage of the overall Edge WAE cache that prepositioned files can consume. For example, if you do not want this prepositioning directive to consume more than 30 percent of a WAE's cache, enter 30 in this field. The default value is 5 percent. The percentage of the cache defined for a preposition task defines the maximum size that can be prepositioned in a single iteration of the task regardless of how much is already in the cache. When the cache is full, regardless of the reason, prepositioning operates like on-demand caching: an eviction process begins and the files with the oldest time-last-accessed values are removed from the cache.
Max File Size	Maximum file size that can be exported. Files that are larger than this value are not exported to the WAE cache.
Min File Size	Minimum file size that can be exported. Files that are smaller than this value are not exported to the WAE cache. It is inefficient to preposition files smaller than 20 KB because these files can be retrieved quickly over the WAN through normal WAAS. The default value is 20 KB.
Duration	Maximum amount of time it should take WAAS to export the file server. If it takes WAAS longer than this amount of time to export the file server, WAAS stops the exporting process before all files are copied to the Edge WAE cache. If the preposition task does not start at the scheduled start time (for example, because the Edge and the Core have no connection), the start retries are counted in the duration. If you do not specify a value for this field, WAAS takes as much time as needed to export this file server.
Type	Time filter on the scan process. From the Type drop-down list, choose one of the following options: • **All Files**—Exports all files to the Edge WAE cache. This is the default setting. • **Files changed since last preposition**—Exports only the files that have changed since the last preposition to the Edge WAE cache. This differential filter is applied from the second iteration of a task execution onward. If a new directory is moved to an already prepositioned directory (without changing its last-modified time), this new directory is not prepositioned during the next prepositioning session when you choose this option. • **Files changed since last**—Exports only the files that have changed within the specified time. For example, if you want to push out file updates that have been made on the file server in the last two hours, enter **2** in the provided field and choose **hour** from the drop-down list.

Note If one of these limits is exceeded during a prepositioning task, the task is terminated and a message is sent to the Administrator log. Any remaining files are exported the next time the task is run. If a user requests one of the missing files before this happens, it is fetched over the WAN through WAAS software as usual.

Step 6 (Optional) Check the **Ignore Hidden Directories** check box if you want to prevent hidden directories on the file server from being prepositioned. This check box is unchecked by default. If you leave this box unchecked, hidden directories are prepositioned.

Step 7 In the File Server field, enter the name of a file server to export. Do not use the double quote (") or forward slash (/) characters.

Step 8 From the Location drop-down list, choose the device location that will provide browsing services for the file server; normally this is the data center WAE. For the best browsing performance, specify a location that is close to the file server. The location is used only for browsing; each edge WAE will retrieve prepositioned files directly from the file server, not from this location. For more information on defining locations, see the "Working with Device Locations" section on page 3-9.

Step 9 In the User name, Password, and Confirm Password fields, enter the username and password credentials for the file server. If the username is in a Windows domain, specify the domain name as part of the User name field, as follows: domain\username.

The access credentials that you enter must allow read access to the prepositioned root directories and to their parent directories.

Step 10 (Optional) Check the **DSCP value for high priority messages** check box if you want to assign a DSCP marking value to the prepositioning traffic. Choose a DSCP value from the drop-down list or enter a number from 0–63 in the text field.

DSCP is a field in an IP packet that enables different levels of service to be assigned to the network traffic. Levels of service are assigned by marking each packet on the network with a DSCP code and associating a corresponding level of service. DSCP is the combination of IP Precedence and Type of Service (ToS) fields. For more information, see RFC 2474.

Step 11 In the Root Share and Directories field, enter the directories on the file server that you want to export. Use any of the following methods to identify a directory:

- Manually enter one or more directory paths in the following format: *protocol*://*server*/*share* or *server**share*. For example, cifs://win12srv/home or win12srv\home. You may enter multiple lines for multiple directories, with each full directory path on its own line. You cannot specify the root directory (/) as a root share.

  When you define multiple root shares, the preposition sequence that is performed for a single root configuration is repeated for each root serially.

- Click the **Browse** button to browse the directories on the file server. To navigate into a directory, click the file folder icon to the left of the directory name. Check the check box next to the directory that you want to export and then click the **Select Directory** button. The browse window allows you to choose multiple directories.

  The browse function operates best when you choose in the Location drop-down list the location of the nearest CIFS accelerator to the file server. If you do not choose a location, the browse request is sent to all devices that have the CIFS accelerator enabled, and the request may time out.

- Check the **Include Sub Directories** check box to include all subdirectories under the specified root directory. If this option is not selected, only the files in the specified root directory are prepositioned and you cannot select subdirectories when you are browsing.

- Narrow the policy definition to a particular type of file by choosing a pattern operator from the File Name drop-down list and entering the text that describes the pattern in the adjacent text box. For example, enter **ends with .doc**. Do not use a space or the following special characters: | : > < " ? * / \

Step 12 Click **Submit**.

The directive is saved and additional tabs appear at the top of the window.

Assigning Edge Devices to a Preposition Directive

After you create a preposition directive, you need to assign Edge WAEs or device groups to the directive. This task determines which Edge WAEs will store preposition content in their cache.

Note Prepositioning includes the ability to configure multiple roots. See the "Creating a New Preposition Directive" section on page 12-12.

To assign an Edge WAE or device group to a preposition directive, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Configure > CIFS File Services > Preposition**.

The Preposition Directives window appears, which lists the preposition directives that exist on the system.

Step 2 Click the **Edit** icon next to the preposition directive that you want to assign to an Edge WAE or device group.

Step 3 Click one of the following tabs at the top of the window:

- **Assign Edge Devices**—Allows you to select one or more Edge WAEs to assign to this directive.
- **Assign Edge Groups**—Allows you to select a device group to assign to this directive.

The Edge Device Assignments window or the Device Groups Assignments window appears, depending on the selected option.

For either view, the assignments window lets you filter your view of the items in the list. Filtering enables you to find items in the list that match the criteria that you set.

Step 4 Choose the Edge WAEs or device groups to assign to this preposition directive by doing either of the following:

- Click in the taskbar to assign all available Edge WAEs or device groups to this directive.
- Click next to the individual Edge WAE or device group that you want to assign to this directive. The icon changes to when selected.

Note If a device or device group is offline (identified by), then you cannot assign that device or group to this directive. The preposition directive, when assigned to a device group, is applied only to connected Edge devices in the assigned device group.

When assigning a CIFS accelerator preposition directive to a device group, the directive is applied only to those devices enabled for CIFS acceleration in the assigned device group.

Step 5 Click **Submit**.

The icon next to each edge device or device group you selected changes to .

Note If the CIFS accelerator is disabled on a WAE, the WAE is removed from any preposition directives to which it is assigned. Also, the preposition directive is removed from the device's running configuration.

Creating a New Preposition Schedule

Once you create a preposition directive and assign WAEs to the directive, we recommend you create a schedule that determines when and how often prepositioning occurs.

For example, you may want to schedule prepositioning to occur at night to minimize the amount of traffic during business hours. Or you may want to schedule prepositioning to occur on a recurring basis if the exported data changes often. This will help ensure that the WAEs assigned to this directive have the latest file updates in their cache.

When a preposition task is scheduled to begin at the same time for multiple Edge WAEs that are located in different timezones, the task will begin on the Edge WAEs based on the Core WAE timezone. If the clocks of the Edge WAE and the Core WAE are not synchronized, the task will not start on time.

To create a preposition schedule, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Configure > CIFS File Services > Preposition**.

The Preposition Directives window appears, which lists the preposition directives that exist on the system.

Step 2 Click the **Edit** icon next to the preposition directive for which you want to create a schedule.

Step 3 Click the **Schedule** tab at the top of the window.

The Creating New Preposition Schedule window appears. By default, no schedule is configured.

Step 4 Choose one of the following scheduling options:

- **Not Scheduled**—Prepositioning is not scheduled at this time.
- **Now**—Prepositioning occurs within a few minutes after you submit this schedule.

  A Now schedule begins again each time you make a change to the preposition directive and click the **Submit** button. A Now schedule also begins again as soon as an edge device that has been reloaded comes back online.
- **Daily**—Prepositioning occurs daily at the defined time.
- **Date**—Prepositioning occurs at the defined time and date.
- **Weekly**—Prepositioning occurs on the selected days of the week at the defined time.
- **Monthly Days**—Prepositioning occurs on the selected days of the month at the defined time.
- **Monthly Weekdays**—Prepositioning occurs on the defined day (as opposed to a defined date) and time during the month. For example, you can schedule prepositioning to occur on the second Tuesday of every month.

Step 5 Specify a start time for the prepositioning task.

The time is expressed in 24-hour format with 00:00 representing midnight. The time refers to the local time of the Edge WAE where the data is to be prepositioned. If there are multiple Edge WAEs in different time zones, the time refers to the local time of the Core WAE.

Note You cannot schedule a start time for the **Now** option.

Step 6 Click **Submit**.

The message Changes Submitted appears at the bottom of the window confirming that your schedule was saved.

Step 7 Verify that the preposition directive completed successfully by checking the preposition status. For more information, see the "Checking the Preposition Status" section on page 12-18.

Checking the Preposition Status

After you create one or more preposition directives, you can check the status of all the preposition tasks to ensure they completed successfully. If a task does not complete successfully, then some of the prepositioned files may have not been successfully copied to the Edge WAE cache.

To check the status of a prepositioning task, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Configure > CIFS File Services > Preposition**.

The Preposition Directives window appears, which lists the preposition directives that exist on the system.

Step 2 Click the **Edit** icon next to the preposition directive for which you want to check.

Step 3 Click the **Preposition Status** tab at the top of the window. The Preposition Status window appears.

This page displays the following information:

- WAE—The name of each Edge WAE that received the prepositioned files in its cache.
- Start Time—The time the preposition task started.
- Duration—The amount of time in took the preposition task to complete.
- Amount Copied—The amount of data copied to the WAE cache (in bytes).
- Status—Whether the preposition task completed successfully.
- Reason—The reason a preposition task failed.

Step 4 Ensure that the Status column shows Completed.

If this column shows a failure, look in the Reason column for an explanation that can help you troubleshoot why the preposition task failed. After resolving the issue, you can schedule the preposition task to run again now, or wait until the scheduled start time and check the status again later.

Starting and Stopping Preposition Tasks

You can start or stop a preposition task from the Device Manager GUI. For more information, see the "Preposition Option" section on page 11-20.

Configuring the SMB Accelerator

Table 12-2 provides an overview of the steps that you must complete to configure the SMB accelerator.

Table 12-4 Checklist for Configuring SMB Accelerator

Task	Additional Information and Instructions
1. Prepare for file services.	Provides the tasks that you need to complete before enabling and configuring file services on your WAAS devices. For more information, see the "Preparing for File Services" section on page 12-7.
2. Enable SMB acceleration.	Enables and configures the SMB accelerator. For more information, see the "Enabling and Disabling the Global Optimization Features" section on page 13-3.
3. (Optional) Identify dynamic shares.	Identifies the dynamic shares on an exported file server. If your file server uses Access Based Enumeration (ABE) to give users different views of the share, you must configure the dynamic shares on the WAAS Central Manager. For more information, see the "Creating Dynamic Shares for the SMB Accelerator" section on page 12-19.

Creating Dynamic Shares for the SMB Accelerator

Many file servers use dynamic shares, which allow multiple users to access the same share but then be automatically mapped to a different directory based on the user's credentials. Dynamic shares are most commonly used on file servers to set up user home directories. For example, a directory named Home can be set up as a dynamic share on a file server so that each user accessing that share is automatically redirected to their own personal directory.

If a file server contains a dynamic share or is using Access Based Enumeration (ABE), you must register that dynamic share with the WAAS Central Manager as described in this section.

Defining a dynamic share in the WAAS Central Manager allows each user to see a different view of the share and allows the operation of ABE if it is configured on the Windows Server.

Note Dynamic share configuration on the WAAS Central Manager overrides any dynamic share configuration set up directly on the WAE device using the CLI.

Before adding a dynamic share, note the following limitations:

- Each dynamic share on a file server must be unique.
- You can use the WAAS Central Manager GUI to define any directory as a dynamic share. However, if a directory is not set up as a dynamic share on the file server, all users will read or write the same content from the same directory and will not be redirected to different directories based on their credentials.

To add a dynamic share for SMB accelerator, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > File Services > SMB Dynamic Shares**.

A list of dynamic shares appears. The Dynamic Shares window shows all the dynamic shares configured. From this window, you can perform the following tasks:

- Edit the configuration of an existing dynamic share by selecting it and clicking the **Edit** taskbar icon.
- Delete the dynamic share by selecting it and clicking the **Delete** taskbar icon.
- Add a new dynamic share definition, as described in the next steps.

Step 3 Click the **Add Dynamic Share** taskbar icon to add a new dynamic share. The Dynamic Share window appears.

Step 4 In the File Server field, enter a valid FQDN or IP address of the file server with the dynamic share.

If you specify the file server name, the WAE resolves it to an IP address.

Step 5 The IP addresses of the registered file servers are displayed in a drop-down list. Choose a file server.

Step 6 In the Share field, specify the location of the dynamic share by doing one of the following tasks:

- Enter the name of the dynamic share on the file server. The following characters cannot be used in the share name: \ / : * ? “ < > |
- Click **Browse** next to the Share Name field to navigate to the correct root directory.

Note The Browse button appears only if you have at least one WAE device with the SMB accelerator enabled and registered to the WAAS Central Manager.

Step 7 Ensure that the status of the share is set to enabled. If you change the status to disabled, the share will not be set up as a dynamic share in your WAAS environment.

Step 8 Click **OK**.

The specified directory now functions as a dynamic share on the WAE.

CHAPTER 13

Configuring Application Acceleration

This chapter describes how to configure the optimization policies on your WAAS system that determine the types of application traffic that is accelerated over your WAN.

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

This chapter contains the following sections:

About Application Acceleration

The WAAS software comes with over 150 predefined optimization policies that determine the type of application traffic your WAAS system optimizes and accelerates. These predefined policies cover the most common type of application traffic on your network. For a list of the predefined policies, see Appendix A, "Predefined Optimization Policy."

Each optimization policy contains the following elements:

- Application definition—Identifies general information about a specific application, such as the application name and whether the WAAS Central Manager collects statistics about this application.
- Class Map—Contains a matching condition that identifies specific types of traffic. For example, the default HTTP class map matches all traffic going to ports 80, 8080, 8000, 8001, and 3128. You can create up to 512 class maps and 1024 matching conditions.
- Policy—Combines the application definition and class map into a single policy. This policy also determines what optimization and acceleration features (if any) a WAAS device applies to the defined traffic. You can create up to 512 policies. A policy can also contain a differentiated services code point (DSCP) marking value that is applied to the traffic and that overrides a DSCP value set at the application or global level.

You can use the WAAS Central Manager GUI to modify the predefined policies and to create additional policies for other applications. For more information on creating optimization policies, see the "Creating a New Traffic Optimization Policy" section on page 13-49. For more information on viewing reports, restoring policies, monitoring applications, and other functions, see the "Managing Application Acceleration" section on page 13-55.

Note All application definitions configured in the WAAS Central Manager are globally applied to all WAAS devices that register with the WAAS Central Manager, regardless of the device group membership configuration.

WAAS policies can apply two kinds of optimizations to matched traffic:

- Layer 4 optimizations that include TFO, DRE, and LZ compression. These features can be applied to all types of TCP traffic.
- Layer 7 optimizations that accelerate application-specific protocols. The application accelerators control these kinds of optimizations.

For a given optimization policy, the DRE feature can use different caching modes (beginning with software version 4.4.1):

- Bidirectional—The peer WAEs maintain identical caches for inbound and outbound traffic. This caching mode is best suited where a significant portion of the traffic seen in one direction between the peers is also seen in the reverse direction. In software versions prior to 4.4.1, this mode is the only supported caching mode.
- Unidirectional—The peer WAEs maintain different caches for inbound and outbound traffic. This caching mode is best suited where a significant portion of the traffic seen in one direction between the peers is not seen in the reverse direction.
- Adaptive—The peer WAEs negotiate either bidirectional or unidirectional caching based on the characteristics of the traffic seen between the peers.

The predefined optimization policies are configured to use the optimal DRE caching mode, depending on the typical application traffic, though you can change the mode if you want.

Enabling and Disabling the Global Optimization Features

The global optimization features determine if TFO Optimization, Data Redundancy Elimination (DRE), and Persistent Compression are enabled on a device or device group. By default, all of these features are enabled. If you choose to disable one of these features, the device will be unable to apply the full WAAS optimization techniques to the traffic that it intercepts.

In addition, the global optimization features include each of the following application accelerators: EPM, CIFS, HTTP, MAPI, NFS, SSL, SMB, ICA, and video. By default, all of the application accelerators are enabled except SMB. Encrypted MAPI is also not enabled by default. The application accelerators also require specific licenses to operate. For information on installing licenses, see the "Managing Software Licenses" section on page 10-3.

You must enable the accelerator on both of the peer WAEs at either end of a WAN link for all application accelerators to operate.

To enable or disable a global optimization feature, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > Enabled Features**.

The Enabled Features window appears. (See Figure 13-1.)

Note For a WAAS Express device, only a subset of the standard features are available. (See Figure 13-2.)

Figure 13-1 Enabled Features Window

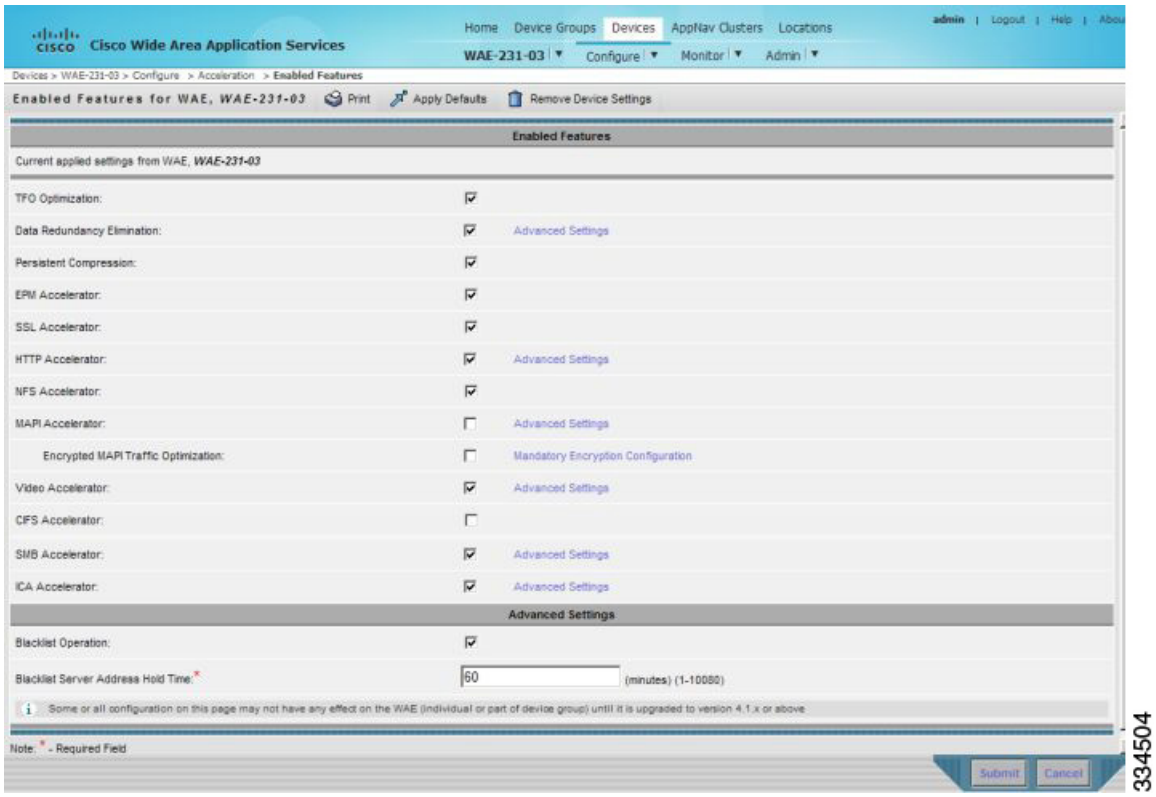

Figure 13-2 shows the subset of standard features that are available for a WAAS Express device.

Figure 13-2 Enabled Features Window—WAAS Express

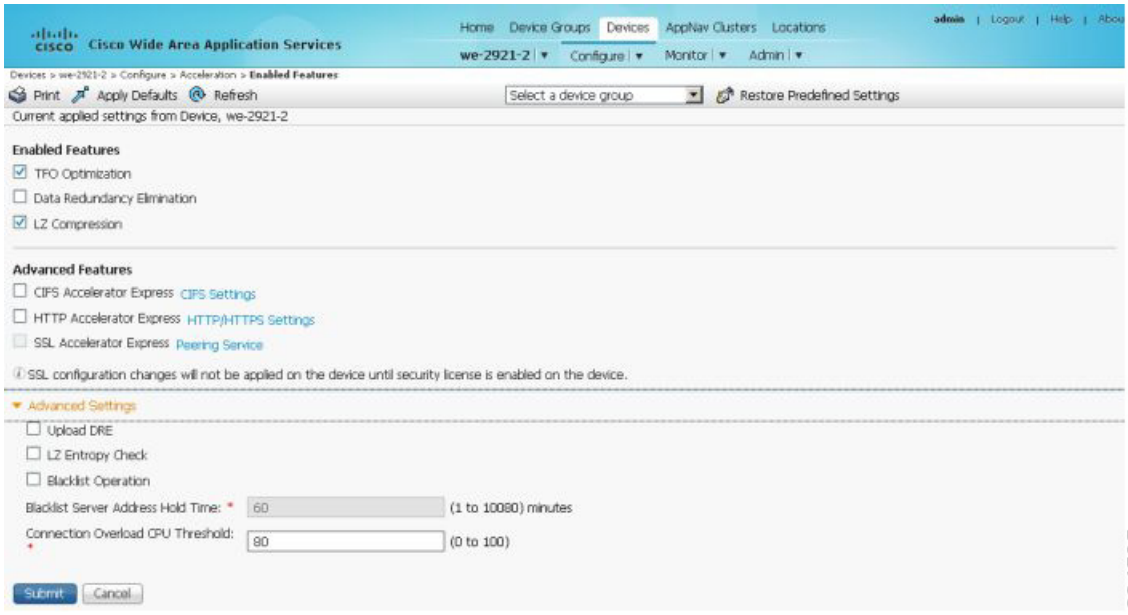

For WAAS Express, the following express versions of application accelerators are supported:

- CIFS accelerator express (See the "Configuring CIFS Accelerator Express" section on page 13-26)
- HTTP accelerator express (See the "Configuring HTTP Acceleration" section on page 13-7)
- SSL accelerator express (See the "Configuring SSL Acceleration" section on page 13-28)

Not all of the properties in the standard WAAS device are available in the WAAS Express version of the application accelerators.

Note If you try to enable DRE on a WAAS Express device on which it is not supported, a message tells you that it is not supported.

The Restore Predefined Settings icon for WAAS Express applies the predefined settings for HTTP/HTTPS, CIFS, and SSL cipher list and peering service.

Step 3 Place a check next to the optimization features that you want to enable, and uncheck the features that you want to disable. For a description of each of the optimization features, see the "Key Services of Cisco WAAS" section on page 1-4.

Step 4 If you check the **Data Redundancy Elimination** check box, you can click the Advanced Settings link as a shortcut to the DRE Settings Configuration window. For more information, see the "Configuring DRE Settings" section on page 13-7.

Step 5 If you check the **HTTP Accelerator** check box, you can click the Advanced Settings link as a shortcut to the HTTP Acceleration Configuration window. For more information, see the "Configuring HTTP Acceleration" section on page 13-7.

Step 6 If you check the **Video Accelerator** check box, you can click the Advanced Settings link as a shortcut to the Video Acceleration Configuration window. For more information, see the "Configuring Video Acceleration" section on page 13-22.

Step 7 If you check the **MAPI Accelerator** check box, you can click the Advanced Settings link as a shortcut to the MAPI Acceleration Configuration window. For more information, see the "Configuring MAPI Acceleration" section on page 13-11.

Step 8 If you check the **Encrypted MAPI Traffic Optimization** check box, you can click the Mandatory Encryption Configuration link as a shortcut to the Encrypted Services Configuration window.

Note The Encrypted MAPI feature is in extended beta trials. You must contact waas-emapi-cs@external.cisco.com with your Cisco account team on the cc: for approvals before enabling this feature. Only approved customers will be supported for beta evaluations.

For more information, see the "Configuring Encrypted MAPI Acceleration" section on page 13-12.

Step 9 If you check the **CIFS Accelerator** check box, you have the following option:

- Windows Print Accelerator—Check this box to accelerate print traffic between clients and a Windows print server. This accelerator is enabled by default when you enable the CIFS accelerator.

Note Do not disable Windows Print Acceleration during a client session as this can interfere with the client's use of print services. If you must disable Windows Print Acceleration, disconnect and then reestablish the client session.

Step 10 If you check the **SMB Accelerator** check box, you can click the Advanced Settings link as a shortcut to the SMB Acceleration Configuration window. For more information, see the "Configuring SMB Acceleration" section on page 13-24.

Step 11 If you check the **ICA Accelerator** check box, you can click the Advanced Settings link as a shortcut to the ICA Acceleration Configuration window. For more information, see the "Configuring ICA Acceleration" section on page 13-27.

Step 12 In the Advanced Settings area, uncheck the **Blacklist Operation** feature if you want to disable it. This feature allows a WAE to better handle situations in which TCP setup packets that have options are blocked or not returned to the WAE device. This behavior can result from network devices (such as firewalls) that block TCP setup packets that have options, and from asymmetric routes. The WAE can keep track of origin servers (such as those behind firewalls) that cannot receive optioned TCP packets and learns not to send out TCP packets with options to these blacklisted servers. WAAS is still able to accelerate traffic between branch and data center WAEs in situations where optioned TCP packets are dropped. We recommend leaving this feature enabled.

Step 13 If you want to change the default Blacklist Server Address Hold Time of 60 minutes, enter the new time in minutes in the Blacklist Server Address Hold Time field. The valid range is 1 minute to 10080 minutes (1 week).

When a server IP address is added to the blacklist, it remains there for configured hold time. After that time, subsequent connection attempts will again include TCP options so that the WAE can redetermine if the server can receive them. It is useful to retry sending TCP options periodically because network packet loss may cause a server to be erroneously blacklisted.

You can shorten or lengthen the blacklist time by changing the Blacklist Server Address Hold Time field.

Step 14 Click **Submit**.

The changes are saved to the device or device group.

To configure TFO optimization, DRE, and persistent compression from the CLI, use the **tfo optimize** global configuration command.

To configure EPM acceleration from the CLI, use the **accelerator epm** global configuration command.

To configure HTTP acceleration from the CLI, use the **accelerator http** global configuration command.

To configure NFS acceleration from the CLI, use the **accelerator nfs** global configuration command.

To configure MAPI acceleration from the CLI, use the **accelerator mapi** global configuration command.

To configure video acceleration from the CLI, use the **accelerator video** global configuration command.

To configure SSL acceleration from the CLI, use the **accelerator ssl** global configuration command.

To configure CIFS acceleration from the CLI, use the **accelerator cifs** and **accelerator cifs preposition** global configuration commands.

To configure Windows print acceleration from the CLI, use the **accelerator windows-print** global configuration command.

To configure SMB acceleration from the CLI, use the **accelerator smb** global configuration command.

To configure ICA acceleration from the CLI, use the **accelerator ica** global configuration command.

To configure the Blacklist Operation feature from the CLI, use the **tfo auto-discovery** global configuration command.

To display status and statistics on the application accelerators from the CLI, use the **show accelerator** and **show statistics accelerator** EXEC commands. To display statistics on the Windows print accelerator, use the **show statistics windows-print requests** EXEC command.

For details on configuring individual application accelerators, see the following sections:

- Configuring HTTP Acceleration, page 13-7

- Configuring MAPI Acceleration, page 13-11
- Configuring Encrypted MAPI Acceleration, page 13-12
- Configuring Video Acceleration, page 13-22
- Configuring CIFS Accelerator Express, page 13-26
- Configuring SMB Acceleration, page 13-24
- Configuring ICA Acceleration, page 13-27
- Configuring SSL Acceleration, page 13-28
- For CIFS: Chapter 12, "Configuring File Services"

Configuring DRE Settings

To enable DRE settings, check the DRE Settings check box in the Enabled Features window (see Figure 13-1 on page 13-4).

To configure the DRE auto bypass and load monitor settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > DRE Settings**.

The DRE Settings window appears.

Step 3 Check the **Enable DRE auto bypass** check box to generate an alarm and automatically DRE bypass application traffic.

Step 4 Check the **Enable DRE Load Monitor** check box to enable load report.

Step 5 Click **Submit**.

The changes are saved to the device or device group.

To enable DRE auto bypass from the CLI, use the **dre auto-bypass enable** global configuration command.

To enable DRE load monitor from the CLI, use the **dre load-monitor report** global configuration command.

Configuring HTTP Acceleration

The HTTP application accelerator accelerates HTTP traffic. SSL traffic that uses HTTPS can be optimized by both SSL and HTTP optimizations.

The default Web optimization policy is defined to send traffic to the HTTP accelerator. The Web optimization policy uses the HTTP class map, which matches traffic on ports 80, 8080, 8000, 8001, and 3128. If you expect HTTP traffic on other ports, add the other ports to the HTTP class map.

To enable the HTTP accelerator, check the HTTP Accelerator check box in the Enabled Features window (see Figure 13-1 on page 13-4).

To configure the HTTP acceleration settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > HTTP/HTTPS Settings**.

The HTTP Acceleration Settings window appears. (See Figure 13-3.)

Note For WAAS Express, the HTTP acceleration settings are the same but the fields are laid out differently in the HTTP/HTTPS Settings window.

Figure 13-3 HTTP Acceleration Settings Window

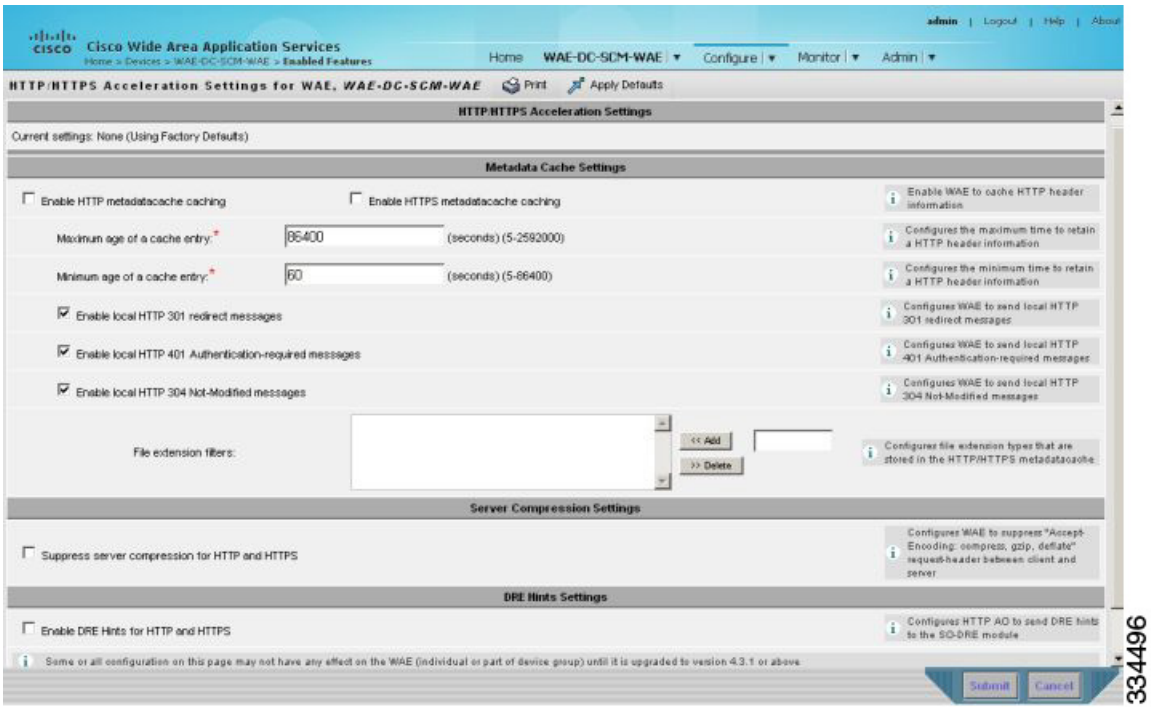

Step 3 Check the **Enable HTTP metadatacache caching** check box to enable the WAE to cache HTTP header (metadata) information. The default setting is enabled.

This box must be checked to enable any of the other settings in the Metadata Cache Settings area. If this box is not checked, no header caching is done.

For details on HTTP metadata caching, see the "About HTTP Metadata Caching" section on page 13-9.

Step 4 Check the **Enable HTTPS metadatacache caching** check box to enable the WAE to cache HTTPS header (metadata) information (HTTP as payload in SSL traffic). The default setting is checked (enabled).

For details on HTTP metadata caching, see the "About HTTP Metadata Caching" section on page 13-9.

Step 5 In the Maximum age of a cache entry field, enter the maximum number of seconds to retain HTTP header information in the cache. The default is 86400 seconds (24 hours). Valid time periods range from 5–2592000 seconds (30 days).

Step 6 In the Minimum age of a cache entry field, enter the minimum number of seconds to retain HTTP header information in the cache. The default is 60 seconds. Valid time periods range from 5 to 86400 seconds (24 hours).

Step 7 Check the **Enable local HTTP 301 redirect messages** check box to enable the WAE to cache and locally serve HTTP 301 messages. The default setting is checked.

Step 8 Check the **Enable local HTTP 401 Authentication-required messages** check box to enable the WAE to cache and locally serve HTTP 401 messages. The default setting is checked.

Step 9 Check the **Enable local HTTP 304 Not-Modified messages** check box to enable the WAE to cache HTTP 200 and 304 messages and locally serve HTTP 304 messages. The default setting is checked.

Step 10 To configure specific file extensions to which metadata caching is to be applied, enter the file extensions in the File extension filters field at the far right. Separate multiple extensions with a comma (for example: jpeg, gif, png) and do not include the dot at the beginning of the file extension. Click the **<< Add** button to add the entered file extensions to the active list, which is shown to the left. You can enter a maximum of 20 file extensions.

To remove an extension from the list, select it in the active list and click the **>> Delete** button.

By default, no file extension filters are defined and therefore metadata caching applies to all file types.

Step 11 Check the **Suppress server compression for HTTP and HTTPS** check box to configure the WAE to suppress server compression between the client and the server. The default setting is checked.

By checking this box, you are telling the WAE to remove the Accept-Encoding value from HTTP and HTTPS request headers, preventing the web server from compressing HTTP and HTTPS data that it sends to the client. This allows the WAE to apply its own compression to the HTTP and HTTPS data, typically resulting in much better compression than the web server for most files. For some file types that rarely change, such as .css and .js files, this setting is ignored and web server compression is allowed.

Step 12 Check the **Enable DRE Hints for HTTP and HTTPS** check box to send DRE hints to the DRE module for improved DRE performance. The DRE hint feature is enabled by default.

Step 13 Click **Submit**.

The changes are saved to the device or device group.

To configure HTTP acceleration from the CLI, use the **accelerator http** global configuration command.

To show the contents of the metadata cache, use the **show cache http-metadatacache** EXEC command.

To clear the metadata cache, use the **clear cache http-metadatacache** EXEC command.

To enable or disable specific HTTP accelerator features for specific clients or IP subnets, use the HTTP accelerator subnet feature. For more details, see the "Using an HTTP Accelerator Subnet" section on page 13-10.

About HTTP Metadata Caching

The metadata caching feature allows the HTTP accelerator in the branch WAE to cache particular server responses and respond locally to clients. The following server response messages are cached:

- HTTP 200 OK (Applies to If-None-Match and If-Modified-Since requests)
- HTTP 301 redirect
- HTTP 304 not modified (Applies to If-None-Match and If-Modified-Since requests)

- HTTP 401 authentication required

Metadata caching is not applied in the following cases:

- Requests and responses that are not compliant with RFC standards
- URLs over 255 characters
- 301 and 401 responses with cookie headers
- HEAD method is used
- Pipelined transactions

Note The metadata caching feature is introduced in WAAS version 4.2.1, but version 4.2.1 is needed only on the branch WAE. This feature can interoperate with an HTTP accelerator on a data center WAE that has a lower version.

Using an HTTP Accelerator Subnet

The HTTP accelerator subnet feature allows you to selectively enable or disable specific HTTP optimization features for specific IP subnets by using ACLs. This feature can be applied to the following HTTP optimizations: HTTP metadata caching, HTTPS metadata caching, DRE hints, and suppress server compression.

To define IP subnets, use the **ip access-list** global configuration command. Refer to this command in the *Cisco Wide Area Application Services Command Reference* for information on configuring subnets. You can use both standard and extended ACLs.

To configure a subnet for an HTTP accelerator feature, follow these steps:

Step 1 Enable the global configuration for all the HTTP accelerator features that you want to use.

Step 2 Create an IP access list to use for a subnet of traffic.

```
WAE(config)# ip access-list extended md_acl
WAE(config-ext-nacl)# permit ip 1.1.1.0 0.0.0.255 any
WAE(config-ext-nacl)# permit ip 2.2.2.0 0.0.0.255 3.3.3.0 0.0.0.255
WAE(config-ext-nacl)# exit
```

Step 3 Associate the ACL with a specific HTTP accelerator feature. Refer to the **accelerator http** global configuration command in the *Cisco Wide Area Application Services Command Reference* for information on associating an ACL with an HTTP accelerator feature.

```
WAE(config)# accelerator http metadatacache access-list md_acl
```

In this example, the HTTP metadata cache feature applies to all the connections that match the conditions specified in the extended access-list md_acl.

In the following example, the HTTP suppress-server-encoding feature applies to all the connections that match the conditions specified in the standard access-list 10.

```
WAE(config)# ip access-list standard 10
WAE(config-std-nacl)# permit 1.1.1.0 0.0.0.255
WAE(config-std-nacl)# exit
WAE(config)# accelerator http suppress-server-encoding accesslist 10
```

For the features (DRE hints and HTTPS metadata cache in this example) that do not have an ACL associated with them, the global configuration is used and they are applicable to all the connections.

Configuring MAPI Acceleration

The MAPI application accelerator accelerates Microsoft Outlook Exchange traffic that uses the Messaging Application Programming Interface (MAPI) protocol. Microsoft Outlook 2000–2010 clients are supported. Clients can be configured with Outlook in cached or noncached mode; both modes are accelerated.

Secure connections that use message authentication (signing) are not accelerated, and MAPI over HTTP is not accelerated.

Note Microsoft Outlook 2007 and 2010 have encryption enabled by default. You must disable encryption to benefit from the MAPI application accelerator.

The EPM application accelerator must be enabled for the MAPI application accelerator to operate. EPM is enabled by default. Additionally, the system must define an optimization policy of type EPM, specify the MAPI UUID, and have an Accelerate setting of MAPI. This policy, MAPI for the Email-and-Messaging application, is defined by default.

EPM traffic, such as MAPI, does not normally use a predefined port. If your Outlook administrator has configured Outlook in a nonstandard way to use a static port, you must create a new basic optimization policy that accelerates MAPI traffic with a class map that matches the static port that was configured for Outlook.

Note If the WAE becomes overloaded with connections, the MAPI application accelerator continues to accelerate MAPI connections by using internally reserved connection resources. If the reserved resources are also exceeded, new MAPI connections are passed through until connection resources become available.

To enable the MAPI accelerator, check the MAPI Accelerator check box in the Enabled Features window (see Figure 13-1 on page 13-4).

To configure MAPI acceleration settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > MAPI Settings**.

The MAPI Acceleration Settings window appears. (See Figure 13-4.)

Figure 13-4 MAPI Acceleration Settings Window

Step 3 In the **Reserved Pool Size Maximum Percent** field, enter the maximum percent of connections to restrict the maximum number of connections reserved for MAPI optimization during TFO overload. It is specified as a percent of the TFO connection limit of the platform. Valid percent ranges from 5%-50%. The default is 15%, which would reserve approximately 0.5 connection for each client-server Association Group (AG) optimized by the MAPI accelerator.

The client maintains at least one AG per server it connects to with an average of about 3 connections per AG. For deployments that observe a greater average number of connections per AG, or where TFO overload is a frequent occurrence, a higher value for reserved pool size maximum percent is recommended.

Reserved connections would remain unused when the device is not under TFO overload. Reserved connections are released when the AG terminates.

Step 4 Click **Submit**. The changes are saved to the device or device group.

Configuring Encrypted MAPI Acceleration

The Encrypted MAPI acceleration feature provides WAN optimization for secure MAPI application protocols using MS-Kerberos security protocol and Windows Active Directory identity for authentication of clients and/or servers in the domain.

Note The Encrypted MAPI feature is in extended beta trials. You must contact waas-emapi-cs@external.cisco.com with your Cisco account team on the cc: for approvals before enabling this feature. Only approved customers will be supported for beta evaluations.

This section contains the following topics:

Task Flow for Configuring Encrypted MAPI

To configure Encrypted MAPI traffic acceleration, complete the tasks listed in Table 13-1. These tasks must be performed on both data center and branch WAEs unless specifically noted as not required (or optional).

Table 13-1 Tasks for Configuring Encrypted MAPI

Task	Additional Information and Instructions
1. Configure DNS Settings.	To configure DNS settings, see the "Configuring the DNS Server" section on page 6-26.
2. Configure NTP Settings.	To synchronize the time with Active Directory, see the "Configuring NTP Settings" section on page 10-5.
3. Verify WAE devices are registered and online with the WAAS Central Manager.	To verify WAE devices are registered and online with the WAAS Central Manager, see the "Devices Window" section on page 17-6.
4. Configure SSL Peering Service.	To configure SSL Peering Service, see the "Configuring SSL Peering Service" section on page 13-43.
5. Verify WAN Secure mode is enabled.	To verify WAN Secure mode is enabled, use the **show accelerator wansecure** EXEC command.
6. Configure windows domain settings and perform domain join. (The domain join function automatically creates the machine account in Active Directory.)	To configure Windows Domain Server Authentication settings, see the "Configuring Windows Domain Server Authentication Settings" section on page 7-17 section. Note that performing a domain join of the WAE is not required on branch WAE devices.
7. Configure domain identities (for machine account and optional user accounts).	To configure a machine account identity, see the "Configuring a Machine Account Identity" section on page 13-15. (Optional) To create a user account and configure a user account identity, see the "Creating and Configuring a User Account" section on page 13-17. Note that configuring domain identities is not required on branch WAE devices.
8. Enable Windows Domain Encrypted Service.	To enable the Windows Domain Encrypted Service, navigate to the **Configure > Security > Windows Domain > Encrypted Services** page and check the Enable Encrypted Service check box.
9. Enable Encrypted MAPI Traffic Optimization.	To enable Encrypted MAPI Traffic, see the "Enabling and Disabling the Global Optimization Features" section on page 13-3.

Configuring Encrypted MAPI Settings

To configure Encrypted MAPI acceleration, follow these steps:

Step 1 Configure DNS Settings.

The WAAS DNS server must be part of the DNS system of Windows Active Directory domains to resolve DNS queries for traffic encryption.

To configure DNS settings, see the "Configuring the DNS Server" section on page 6-26.

Step 2 Configure NTP Settings to synchronize the time with Active Directory.

The WAAS device has to be in synchronization with Active Directory for Encrypted MAPI acceleration. The WAAS NTP server must share time synchronization with the Active Directory Domain Controllers domains for which traffic encryption is desired. Out of sync time will cause Encrypted MAPI acceleration to fail.

To synchronize the time with Active Directory, see the "Configuring NTP Settings" section on page 10-5.

Step 3 Verify WAE devices are registered and online with the WAAS Central Manager.

To verify WAE devices are registered and online with the WAAS Central Manager, see the "Devices Window" section on page 17-6.

Step 4 Configure SSL Peering Service.

Note SSL accelerator must be enabled and in the running state.

To configure SSL Peering Service, see the "Configuring SSL Peering Service" section on page 13-43.

Step 5 Verify WAN Secure mode is enabled.

The default mode is Auto. You can verify the state of WAN Secure mode using the following EXEC command:

show accelerator wansecure

If necessary, you can change the state of WAN Secure using the following global configuration command:

accelerator mapi wansecure-mode {**always** | **auto** | **none**}

Step 6 Configure windows domain settings and perform domain join. (Domain join automatically creates the machine account in Active Directory.)

Note Performing a domain join of the WAE is not required on branch WAE devices.

Note This step is optional on data center WAEs if only user accounts are used for domain identity configuration in the next step.

To configure Windows Domain Server Authentication settings, see the "Configuring Windows Domain Server Authentication Settings" section on page 7-17 section.

Note You must use Kerberos authentication for Encrypted MAPI Acceleration. NTLM authentication method is not supported.

Step 7 Configure domain identities. (Not required for branch WAEs.)

You must have at least one account configured, either user or machine, that is configured with a domain identity. Each device can support up to 5 domain identities,1 machine account identity and 4 user account identities. This allows a WAAS device to accelerate up to 5 domain trees. You must configure a domain identity for each domain with an exchange server that has clients to be accelerated.

a. Configure the machine account identity.

A machine account for the core device was automatically created during the join process in the Windows Domain Server authentication procedure in the previous step. If you are using a machine account, a machine account identity must be configured for this account.

Each device only supports one machine account identity.

To configure a machine account identity, see the "Configuring a Machine Account Identity" section on page 13-15.

b. Create and configure optional user accounts.

You may utilize up to four optional user accounts for additional security. Multiple user accounts provide greater security than having all of the core devices using a single user account. You are required to configure a user account identity for each user account, whether you are utilizing an existing user account or creating a new one.

To create a user account and configure a user account identity, see the "Creating and Configuring a User Account" section on page 13-17.

Step 8 Enable Windows Domain Encrypted Service. (Enabled by default.)

a. From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

b. From the menu, choose **Configure > Security > Windows Domain > Encrypted Services**. The Encrypted Services window appears.

c. Check the **Enable Encrypted Service** check box.

d. Click **Submit** to save your changes.

Step 9 Enable Encrypted MAPI Traffic Optimization.

From the Enabled Features window, check the **Encrypted MAPI Traffic Optimization** check box (the **MAPI Accelerator** check box must also be checked), and click **Submit**. Encrypted MAPI traffic optimization is disabled by default.

Note The Encrypted MAPI feature is in extended beta trials. You must contact waas-emapi-cs@external.cisco.com with your Cisco account team on the cc: for approvals before enabling this feature. Only approved customers will be supported for beta evaluations.

For more information on the Enabled Features window, see the "Enabling and Disabling the Global Optimization Features" section on page 13-3.

Configuring a Machine Account Identity

To configure an identity for a machine account, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 From the menu, choose **Configure > Security > Windows Domain > Encrypted Services**. The Encrypted Services window appears.

Step 3 Click the **Add Domain Identity** button to add a machine account domain identity. (See Figure 13-6.)

Every WAAS device to be accelerated must have a domain identity.

Figure 13-5 Add Domain Identity—Machine Account

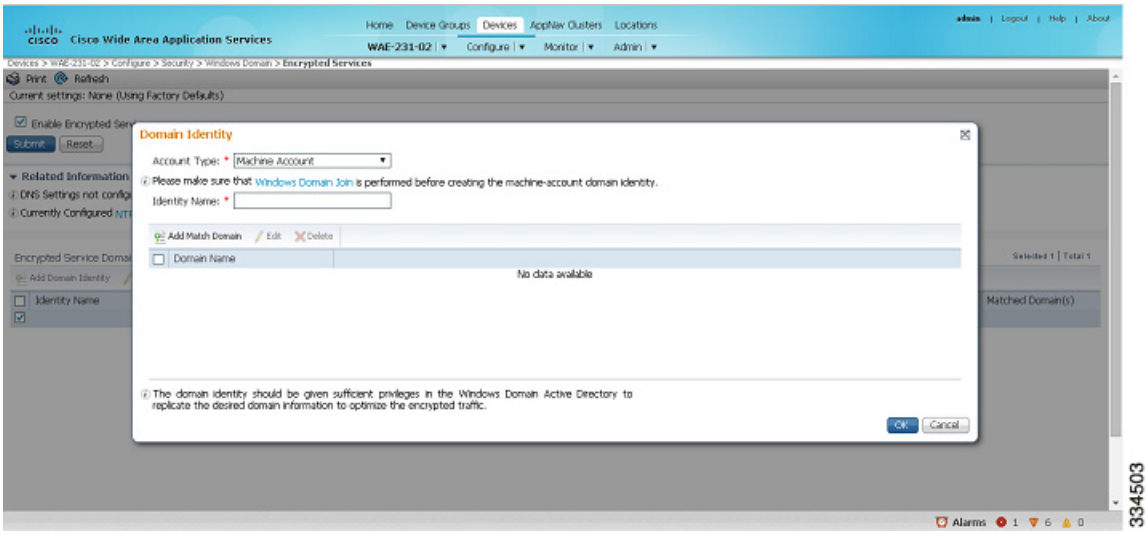

a. Select Machine Account from the **Account Type** drop-down list.

Note Windows Domain Join must be completed before creating the machine account domain identity.

b. Enter the identity name in the **Identity Name** field. Alphanumeric characters only (cannot contain space, ?, |), not exceeding 32 characters.

Note The domain identity must have sufficient privileges in the Windows Domain Active Directory to replicate the desired domain information to optimize encrypted traffic. To configure privileges, see the "Configuring Microsoft Active Directory" section on page 13-18.

Step 4 Click **OK**. The domain identity appears in the Encrypted Services Domain Identities list. (See Figure 13-6.)

Figure 13-6 Encrypted Services—Domain Identity

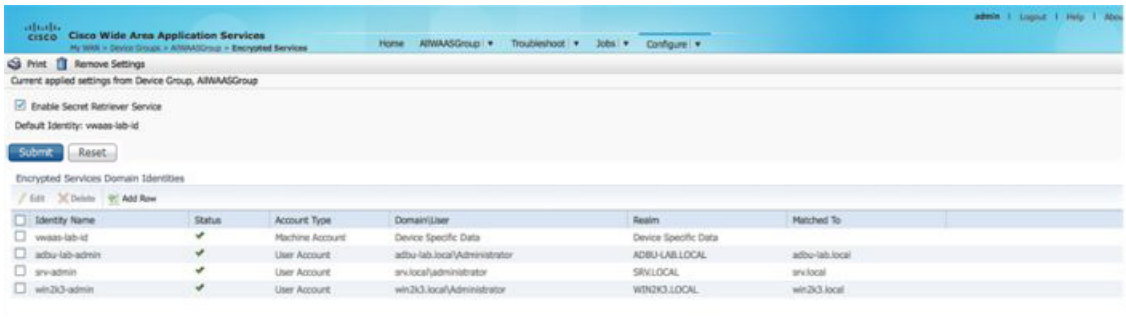

To configure and verify Encrypted Services Domain Identities from the CLI, use the **windows-domain encrypted-service** global configuration command and the **show windows-domain encrypted-service** EXEC command.

Creating and Configuring a User Account

To create a user account and configure a user account identity, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 From the menu, choose **Configure > Security > Windows Domain > Encrypted Services**.

The Encrypted Services window appears. (See Figure 13-7.)

Figure 13-7 Encrypted Services

Step 3 Click the **Add Domain Identity** button to add a user account domain identity. The Domain Identity window appears. (See Figure 13-8.)

Figure 13-8 Add Domain Identity—User Account

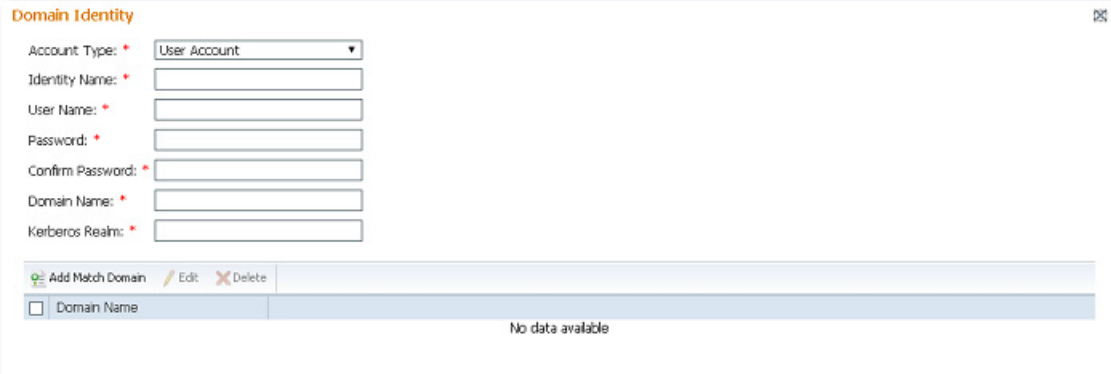

a. Select user account from the **Account Type** drop-down list.

b. Enter the identity name in the **Identity Name** field. Alphanumeric characters only (cannot contain space, ?, |), not exceeding 32 characters.

c. Enter username and password information.

d. Enter the domain name.

e. Enter the Kerberos realm.

Note The domain identity must have sufficient privileges in the Windows Domain Active Directory to replicate the desired domain information to optimize encrypted traffic. To configure privileges, see the "Configuring Microsoft Active Directory" section on page 13-18.

Step 4 Click **OK**. The domain identity appears in the Encrypted Services Domain Identities list.

To configure and verify Encrypted Services Domain Identities from the CLI, use the **windows-domain encrypted-service** global configuration command and the **show windows-domain encrypted-service** EXEC command.

Configuring Microsoft Active Directory

To grant Cisco WAAS permission to accelerate Exchange encrypted email sessions, follow these steps:

Step 1 Using an account with Domain Administrator privileges, launch the Active Directory Users and Computers application.

Step 2 Create a new group.

Note This group is for accounts that WAAS will use to optimize Exchange traffic. Normal users and computers should not be added to this group.

a. Right-click the Organizational Unit (OU) to contain the new group and choose **New > Group**.(See Figure 13-9.)

Figure 13-9 Active Directory—Add Group

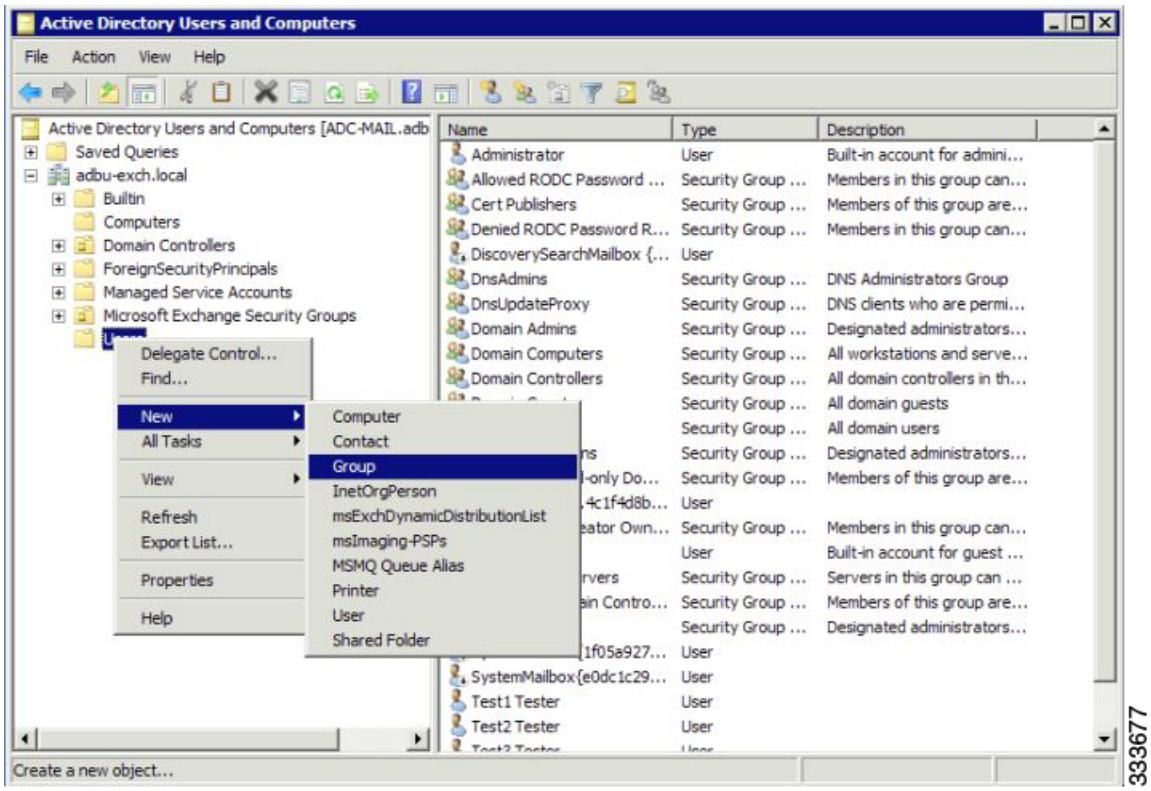

b. Enter a name in the **Group name** fields and select the following attributes:

- Group scope: Universal
- Group type: Security

c. Click **OK**.

Step 3 Configure the permissions required by WAAS.

a. In the Active Directory Users and Computers application window, select **View > Advanced Features** from the menu bar.

b. Right-click on the root of the domain and choose **Properties**.

c. Select the **Security** tab. (See Figure 13-10.)

Figure 13-10 Active Directory—Security Tab

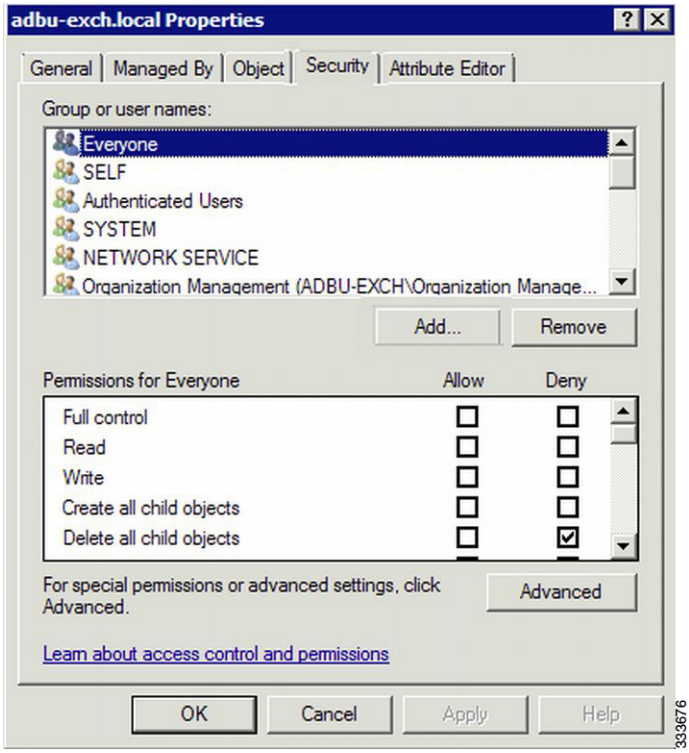

d. Click **Add** in the Group or User Names section.

e. Enter the name of the new group that you created in this procedure in the Enter the object names to select field and click **OK** to add the new group to the list.

f. Select the new group in the Group or user names list and set the following permissions to **Allow**:

- Replicating Directory Changes
- Replicating Directory Changes All

g. Click **OK**.

Step 4 Add an account to the group.

User or workstation (computer) accounts must be added to the new group for WAAS Exchange Encrypted email optimization.

a. Right-click on the account you want to add and select the **Member Of** tab.

b. Click **Add**.

c. Choose the new group you created and click **OK**.

Active Directory permissions configuration is complete.

Managing Domain Identities and Encrypted MAPI State

This section contains the following topics:

- Editing an Existing Domain Identity, page 13-21

Editing an Existing Domain Identity

You can modify the attributes of an existing domain identity on a WAAS device, if needed.

Note If the password for a user account has been changed in Active Directory, you must edit the user account domain identity on the WAAS device to match the new Active Directory password.

The following restrictions apply:

- For a machine account identity, only the state of the domain identity (enabled or disabled) can be modified from a WAAS device.
- For a user account identity, only the state of the domain identity (enabled or disabled) and the password can be modified from a WAAS device.

To change the password for a user account domain identity on a WAAS device when the password for the account in Active Directory has changed, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 From the menu, choose **Configure > Security > Windows Domain > Encrypted Services**.

The Encrypted Services window appears.

Step 3 Select the user account domain identity to modify and click the **Edit** icon.

The Domain Identity window appears.

Step 4 Change the password in the **password** field. The password should be the same as the password for the account in Active Directory.

Step 5 Click **OK**.

Deleting an Existing Domain Identity

To delete a domain identity on a WAAS device, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 From the menu, choose **Configure > Security > Windows Domain > Encrypted Services**.

The Encrypted Services window appears.

Step 3 Select one or more domain identities to delete and click the **Delete** icon to remove the domain identity configured on the WAAS device.

A warning message appears if the domain identity is being used for optimizing encrypted traffic.

Step 4 Click **OK** to accept or **Cancel** to abort the procedure.

Disabling Encrypted MAPI

To disable Encrypted MAPI, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Disable Encrypted Service.

a. From the menu, choose **Configure > Security > Windows Domain > Encrypted Services**.

The Encrypted Services window appears.

b. Uncheck the **Enable Encrypted Service** check box.

c. Click **Submit** to save your changes.

Step 3 Disable Encrypted MAPI Traffic Optimization.

a. From the menu, choose **Configure > Acceleration > Enabled Features**.

The Enabled Features window appears.

b. Uncheck the **Encrypted MAPI Traffic Optimization** check box.

c. Click **Submit** to save your changes.

Encrypted MAPI Acceleration Statistics

To view statistics for Encrypted MAPI connections, see the "Using Predefined Reports to Monitor WAAS" section on page 17-35 and see the MAPI acceleration reports.

Configuring Video Acceleration

The video application accelerator accelerates Windows Media live video broadcasts that use RTSP over TCP. The video accelerator automatically splits one source video stream from the WAN into multiple streams to serve multiple clients on the LAN.

The video accelerator automatically causes the client that is requesting a UDP stream to do a protocol rollover to use TCP (if both the client and server allow TCP).

The default RTSP class map for the Streaming optimization policy is defined to send traffic to the video accelerator.

By default, the video accelerator sends any unaccelerated video traffic to be handled by the negotiated standard TCP optimization policy unless the video accelerator is explicitly configured to drop such traffic. You can choose to drop all unaccelerated video traffic or only traffic that is unaccelerated due to an overload condition.

To enable the video accelerator, check the Video Accelerator check box in the Enabled Features window (see Figure 13-1 on page 13-4).

To configure the video acceleration settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > Video Settings**.

The Video Acceleration Configuration window appears. (See Figure 13-11.)

Figure 13-11 Video Acceleration Configuration Window

Step 3 In the Client First Message Reception Timeout field, enter the number of seconds to wait for the first message from the client and the first response from the server, after the connection is accepted by the video accelerator, before timing out the connection. Valid values range from 10–180 seconds. The default is 60.

Step 4 In the drop-down list, choose which unaccelerated video traffic to drop, as follows:

- **All**—Drop all video traffic that is not being accelerated due to an unsupported transport or format, or overload. All Windows Media video-on-demand traffic and all non-Windows Media RTSP traffic is dropped.
- **Overload Only**—Drop all video traffic that is not being accelerated due to an accelerator overload only.
- **None**—Handle unaccelerated video connections with the negotiated TCP optimization policy. (The traffic is not dropped.)

Note Under some conditions, the video accelerator is not registered with the policy engine, such as when there is no valid license or in certain error conditions. If you configure the video accelerator to drop all unaccelerated video traffic, the policy engine drops all video traffic (even traffic that would have been accelerated if the video accelerator had been properly registered with the policy engine).

Step 5 Check the **Enable transaction logs** check box to enable transaction logging. This feature will generate a large amount of logging data. This box is unchecked by default. Click the **More Settings** link to go to the Windows Media Transaction Log Settings configuration page.

Step 6 Check the **Enable log forwarding** check box to enable forwarding of Windows Media logs to the upstream Windows Media Server. This box is checked by default.

Step 7 In the **Client Idle Connection timeout** field, enter the maximum number of seconds to wait after the initial client request, while the client connection is idle, before timing out the connection. Valid values range from 30–300 seconds. The default is 60.

Step 8 Click **Submit**.

The changes are saved to the device or device group.

To configure video acceleration from the CLI, use the **accelerator video** global configuration command.

Configuring SMB Acceleration

The SMB application accelerator handles optimizations of file server operations. It can be configured to perform the following file server optimizations:

- Read Ahead optimization—The SMB accelerator performs a read-ahead optimization on files that use the oplock feature. When a client sends a read request for a file, it is likely that it may issue more read requests for the same file. To reduce the use of network bandwidth to perform these functions over the WAN on the file server, the SMB accelerator performs read-ahead optimization by proactively reading more file data than what has been initially requested by the client.
- Directory listing optimization—A significant portion of the traffic on the network is for retrieving directory listings. The SMB accelerator optimizes directory listings from the file server through prefetching. For directory prefetching, a request from the client is expanded to prefetch up to 64 KB of directory listing content. The SMB accelerator buffers the pre-fetched directory listing data until the client has requested all the data. If the directory listing size exceeds 64 KB then a subsequent request from the client is expanded by the SMB accelerator again to prefetch content up to 64 KB. This continues until all the entries of the directory are returned to the client.
- Metadata optimization—The SMB accelerator optimizes fetching metadata from the file server through metadata prefetching. Additional metadata requests are tagged along with the client request and are sent to the file server to prefetch more information levels than what was requested by the client.
- Named Pipe optimization—The SMB accelerator optimizes frequent requests from Windows Explorer to the file server to retrieve share, server, and workstation information. Each of these requests involves a sequence of operations that include opening and binding to the named pipe, making the RPC request, and closing the named pipe. Each operation incurs a round trip to the file server. To reduce the use of network bandwidth to perform these functions over the WAN on the file server, the SMB accelerator optimizes the traffic on the network by caching named pipe sessions and positive RPC responses.
- Write optimization—The SMB accelerator performs write optimization by speeding up the write responses to the client by acknowledging the Write requests to the client whenever possible and, at the same time, streaming the Write request over the WAN to the server.
- Not-Found Metadata caching—Applications sometimes send requests for directories and files that do not exist on file servers. For example, Windows Explorer accesses the Alternate Data Streams (ADS) of the file it finds. With negative Not-Found (NF) metadata caching, the full paths to those nonexistent directories and files are cached so that further requests for the same directories and files get local denies to save the round-trips of sending these requests to the file servers.
- DRE-LZ Hints—The SMB accelerator provides DRE hints to improve system performance and resources utilization. At the connection level, the SMB accelerator uses the BEST_COMP latency sensitivity level for all connections, as it gives the best compression. At the message level, the SMB accelerator provides message-based DRE hints for each message to be transmitted over the WAN.
- Microsoft optimization—The SMB accelerator optimizes file operations for Microsoft applications by identifying lock request sequences for file name patterns supported by Microsoft Office applications.

- Invalid FID optimization—The SMB accelerator optimizes SMB2 clients by locally denying attempts to access files with invalid file handle values instead of sending such requests to the file servers.
- Batch Close optimization—The SMB accelerator performs asynchronous file close optimizations on SMB2 traffic.

To enable the SMB accelerator, check the SMB Accelerator check box in the Enabled Features window.

Note The CIFS accelerator and SMB accelerator are mutually exclusive. Both of these cannot be enabled at the same time.

To configure the SMB acceleration settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure** > **Acceleration** > **SMB Settings**.

The SMB Optimization Bypass Settings window appears. (See Figure 13-12.)

Figure 13-12 SMB Accelerator Configuration Window

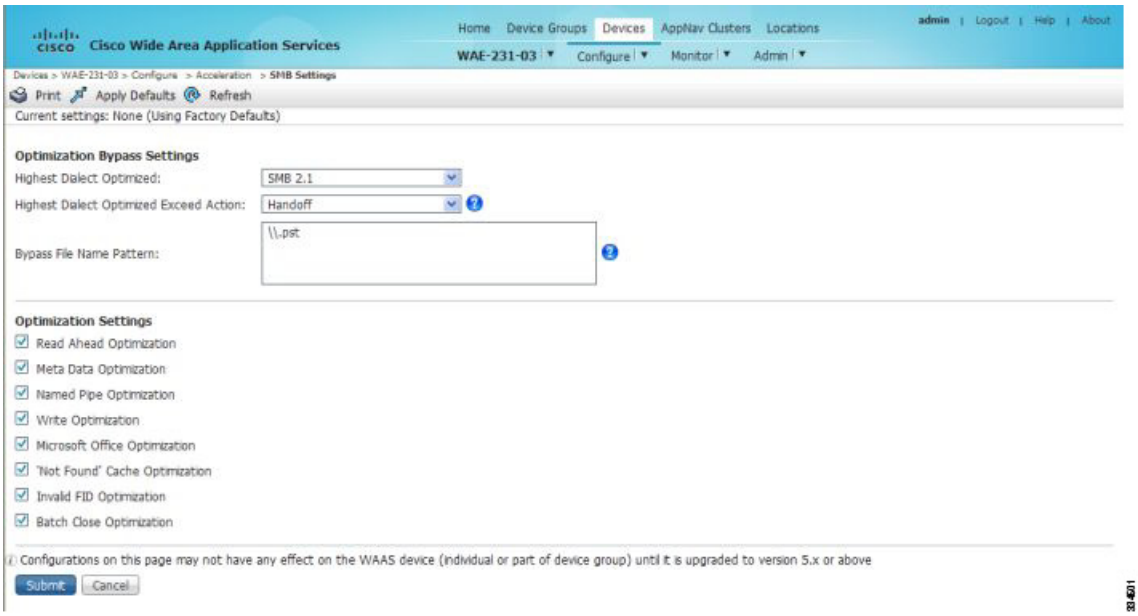

Step 3 In the Highest Dialect Optimized drop-down list, choose the highest dialect to optimize. The available options are:

- NTLM 0.12 or NTLM 1.0
- SMB 2.0
- SMB 2.1

Step 4 In the Highest Dialect Optimized Exceed Action drop-down list, choose the action for the dialects that are higher than the one chosen as the highest dialect to optimize, as follows:

- **Handoff**—If the negotiated dialect is higher than the chosen highest dialect to optimize, the connection is handed off to the generic accelerator.
- **Mute**—The dialects higher than the one chosen as the highest dialect to optimize are removed from the negotiation list.

Step 5 In the Bypass File Name Pattern text box, enter the patterns for the file names that you want the SMB accelerator to bypass optimization for. The files whose names match the specified expressions are not optimized.

Step 6 Check the **Read Ahead Optimization** check box to enable the SMB to optimize the quantity of read-ahead data from the file. The SMB performs a read-ahead optimization only when the file is opened using the ops lock feature. This box is checked by default.

Step 7 Check the **Meta Data Optimization** check box to enable metadata optimization. This box is checked by default.

Step 8 Check the **Named Pipe Optimization** check box to enable named pipe optimization by caching named pipe sessions and positive RPS responses. This box is checked by default.

Step 9 Check the **Write Optimization** check box to enable the write optimization by speeding up the write responses to the client. This box is checked by default

Step 10 Check the **Microsoft Office Optimization** check box to enable optimizations for all versions of Microsoft Office. The SMB accelerator does not perform read-ahead, write optimization, and lock-ahead for Microsoft Office if this optimization is disabled. This box is checked by default.

Step 11 Check the **'Not Found' Cache Optimization** check box to enable caching pathnames of files not found. This box is checked by default.

Step 12 Check the **Invalid FID Optimization** check box to enable optimization of handling files with invalid file handle values. This box is checked by default.

Step 13 Check the **Batch Close Optimization** check box to enable asynchronous file close optimizations. This box is checked by default.

Step 14 Click **Submit** to save the changes.

To configure SMB acceleration from the CLI, use the **accelerator smb** global configuration command.

Configuring CIFS Accelerator Express

The CIFS application accelerator express handles optimizations of file server operations on a WAAS Express device. It interoperates with either the standard CIFS accelerator or the standard SMB accelerator on a standard WAAS device.

CIFS accelerator express can be configured to perform the following file server optimizations:

- Write optimization—CIFS accelerator express performs write optimization by speeding up the write responses to the client by acknowledging the Write requests to the client whenever possible and, at the same time, streaming the Write request over the WAN to the server.
- Read Ahead optimization—CIFS accelerator express performs a read-ahead optimization on files that use the oplock feature. When a client sends a read request for a file, it is likely that it may issue more read requests for the same file. To reduce the use of network bandwidth to perform these functions over the WAN on the file server, the SMB accelerator performs read-ahead optimization by proactively reading more file data than what has been initially requested by the client.

- ADS Negative Cache—Applications sometimes send requests for directories and files that do not exist on file servers. For example, Windows Explorer accesses the Alternate Data Streams (ADS) of the file it finds. With ADS Negative caching, the full paths to those nonexistent directories and files are cached so that further requests for the same directories and files get local denies to save the round-trips of sending these requests to the file servers.

To enable CIFS accelerator express, check the CIFS Accelerator Express check box in the Enabled Features window.

To configure the CIFS accelerator express settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure** > **Acceleration** > **CIFS Settings**.

The CIFS Optimization Bypass Settings window appears.

Step 3 Check the **Write Optimization** check box to enable the write optimization by speeding up the write responses to the client. This box is checked by default

Step 4 Check the **Read Ahead Optimization** check box to enable CIFS accelerator express to optimize the quantity of read-ahead data from the file. CIFS accelerator express performs a read-ahead optimization only when the file is opened using the ops lock feature. This box is checked by default.

Step 5 Check the **ADS Negative Cache** check box to enable caching pathnames of files not found. This box is checked by default.

Step 6 Click **Submit** to save the changes.

To configure CIFS accelerator express from the CLI, use the **accelerator cifs** global configuration command.

Configuring ICA Acceleration

The ICA application accelerator provides WAN optimization on a WAAS device for ICA (Independent Computing Architecture) traffic which is used to access a virtual desktop infrastructure (VDI). This is done through a process that is both automatic and transparent to the client and server.

ICA acceleration is enabled on a WAAS device by default.

To enable the ICA accelerator, check the ICA Accelerator check box in the Enabled Features window (see Figure 13-13 on page 13-28).

To configure the ICA acceleration settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > ICA Settings**.

The ICA Acceleration Configuration window appears. (See Figure 13-11.)

Figure 13-13 ICA Acceleration Configuration Window

Step 3 In the WAN Secure Mode drop-down list, choose the mode, as follows:

- **None**—Disables WAN Secure mode for ICA.
- **Always**—Enables WAN Secure mode ICA. This is the default.

Note The state of WAN Secure mode in both Branch WAE and Data Center WAE must match for connections to get optimized with the ICA accelerator.

Step 4 Click **Submit**.

The changes are saved to the device or device group.

To configure ICA acceleration from the CLI, use the **accelerator ica** global configuration command.

To verify the status of WAN Secure mode from the CLI, use the **show accelerator wansecure** EXEC command.

Configuring SSL Acceleration

The SSL application accelerator optimizes traffic on Secure Sockets Layer (SSL) encrypted connections. If SSL acceleration is not enabled, the WAAS software DRE optimizations are not very effective on SSL encrypted traffic. The SSL application acceleration enables WAAS to decrypt and apply optimizations while maintaining the security of the connection.

Note On a WAAS Express device, only SSL cipher list, SSL certificate authorities, and SSL peering service configuration is supported.

Note The SSL accelerator does not optimize protocols that do not start their SSL/TLS handshake from the very first byte. The only exception is HTTPS going through a proxy (where the HTTP accelerator detects the start of SSL/TLS). In this case, both HTTP and SSL accelerators optimize the connection.

The SSL application accelerator supports SSL Version 3 (SSLv3) and Transport Layer Security Version 1 (TLSv1) protocols. TLSv1.1 and TLSv1.2 protocols are not supported.

Table 13-2 provides an overview of the steps you must complete to set up and enable SSL acceleration.

Table 13-2 Checklist for Configuring SSL Acceleration

Task	Additional Information and Instructions
1. Prepare for configuring SSL acceleration.	Identifies the information that you need to gather before configuring SSL acceleration on your WAAS devices. For more information, see the "Preparing to Use SSL Acceleration" section on page 13-29.
2. Enable secure store, the Enterprise License, and SSL acceleration.	Describes how to set up Central Manager secure store, how to enable the Enterprise License, and how to enable SSL acceleration. Secure store mode is required for secure handling of the SSL encryption certificates and keys. For more information, see the "Enabling Secure Store, the Enterprise License, and SSL Acceleration" section on page 13-30.
3. Enable SSL application optimization.	Describes how to activate the SSL acceleration feature. For more information, see the "Enabling and Disabling the Global Optimization Features" section on page 13-3.
4. Configure SSL acceleration settings.	(Optional) Describes how to configure the basic setup of SSL acceleration. For more information, see the "Configuring SSL Global Settings" section on page 13-31.
5. Create and manage cipher lists.	(Optional) Describes how to select and set up the cryptographic algorithms used on your WAAS devices. For more information, see the "Working with Cipher Lists" section on page 13-35.
6. Set up CA certificates.	(Optional) Describes how to select, import, and manage certificate authority (CA) certificates. For more information, see the "Working with Certificate Authorities" section on page 13-37.
7. Configure SSL management services.	(Optional) Describes how to configure the SSL connections used between the Central Manager and WAE devices. For more information, see the "Configuring SSL Management Services" section on page 13-41.
8. Configure SSL peering service.	(Optional) Describes how to configure the SSL connections used between peer WAE devices for carrying optimized SSL traffic. For more information, see the "Configuring SSL Peering Service" section on page 13-43.
9. Configure and enable SSL accelerated services.	Describes how to add, configure, and enable services to be accelerated by the SSL application optimization feature. For more information, see the "Using SSL Accelerated Services" section on page 13-45.

Preparing to Use SSL Acceleration

Before you configure SSL acceleration, you should know the following information:

- The services that you want to be accelerated on the SSL traffic
- The server IP address and port information
- The public key infrastructure (PKI) certificate and private key information, including the certificate common name and certificate authority signing information
- The cipher suites supported
- The SSL versions supported

Figure 13-14 shows how the WAAS software handles SSL application optimization.

Figure 13-14 SSL Acceleration Block Diagram

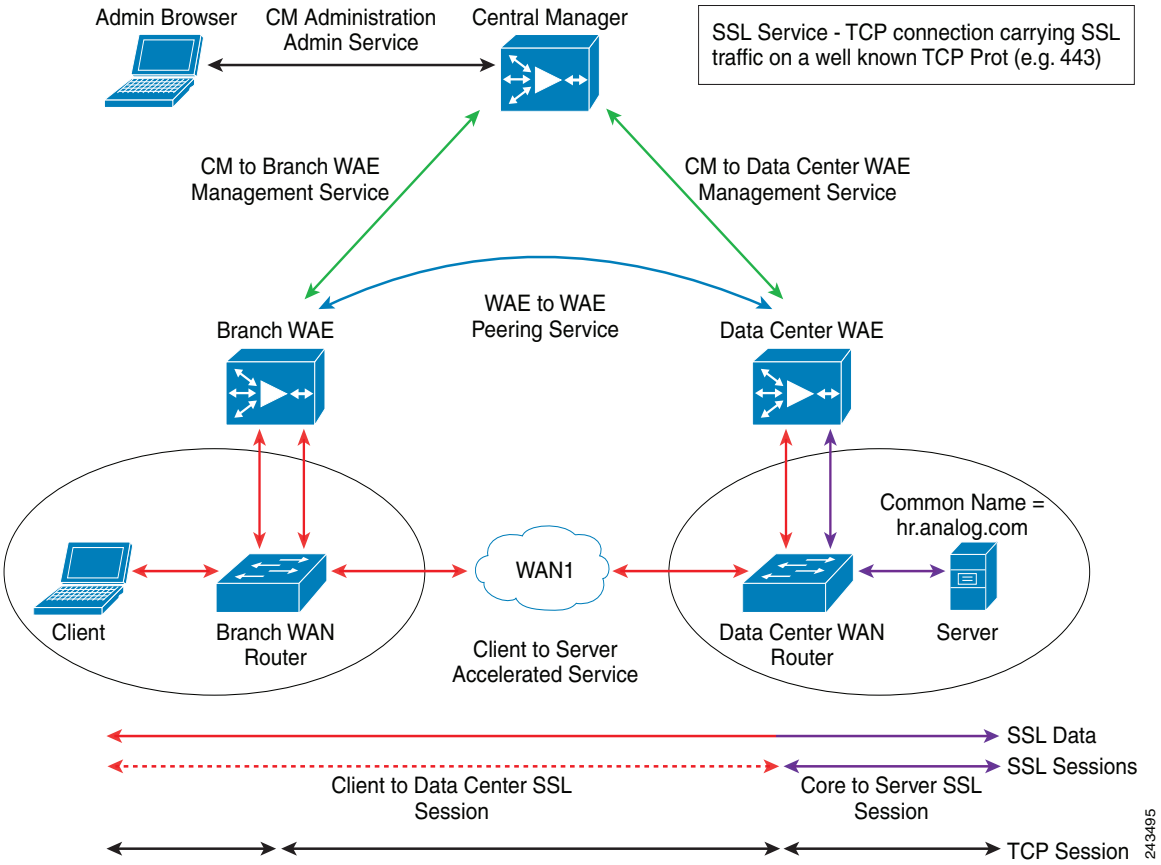

When you configure SSL acceleration, you must configure SSL accelerated service on the server-side (Data Center) WAE devices. The client-side (Branch) WAE needs to have its secure store initialized and unlocked/opened, but does not need to have the SSL accelerated service configured. However, the SSL accelerator must be enabled on both Data Center and Branch WAEs for SSL acceleration services to work. The WAAS Central Manager provides SSL management services and maintains the encryption certificates and keys.

Enabling Secure Store, the Enterprise License, and SSL Acceleration

Before you can use SSL acceleration on your WAAS system, you must perform the following steps:

Step 1 Enable secure store encryption on the Central Manager.

To enable secure store encryption, see the "Configuring Secure Store Settings" section on page 10-10.

Step 2 Enable the Enterprise license.

To enable the Enterprise license, see the "Managing Software Licenses" section on page 10-3.

Step 3 Enable SSL acceleration on devices.

To enable the SSL acceleration feature, see the "Enabling and Disabling the Global Optimization Features" section on page 13-3.

Note If the SSL accelerator is already running, you must wait 2 datafeed poll cycles when registering a new WAE with a Central Manager before making any configuration changes, otherwise the changes may not take effect.

Configuring SSL Global Settings

To configure the basic SSL acceleration settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Security > SSL > Global Settings**.

The SSL Global Settings window appears (see Figure 13-15).

Figure 13-15 SSL Global Settings Window

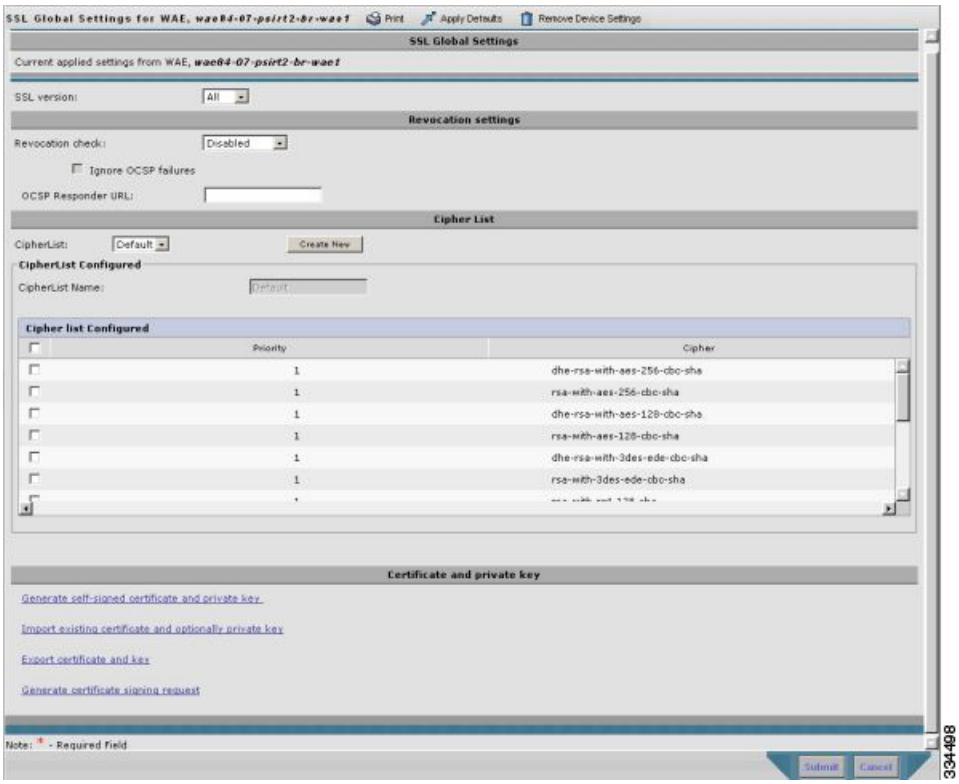

Step 3 To configure a device to use the SSL settings from a particular device group, choose the device group from **Select a Device Group** drop-down list located in SSL global settings toolbar. A device can either use its own SSL settings, or SSL settings from a device group. However, it is not possible to configure a device to use SSL settings from multiple device groups.

Step 4 In the SSL version field, choose the type of SSL protocol to use. Choose **SSL3** for the SSL version 3 protocol, choose **TLS1** for the Transport Layer Security version 1 protocol, or choose **All** to accept both SSL3 and TLS1 SSL protocols.

Step 5 (Optional) Set the Online Certificate Status Protocol (OCSP) parameters for certificate revocation:

a. In the OCSP Revocation check drop-down list, select the OCSP revocation method.

Choose **ocsp-url** SSL accelerator to use OCSP responder specified in the **OCSP Responder URL** field to check the revocation status of certificates. Choose **ocsp-cert-url** to use the OCSP responder URL specified in the Certificate Authority certificate that signed the certificate.

b. If the **Ignore OCSP failures** check box is enabled, the SSL accelerator will treat the OCSP revocation check as successful if it did not get a definite response from the OCSP responder.

Step 6 In the Cipher List field, choose a list of cipher suites to be used for SSL acceleration. For more information, see the "Working with Cipher Lists" section on page 13-35.

Step 7 Choose a certificate/key pair method (see Figure 13-16).

Figure 13-16 Configuring Service Certificate and Private Key

- Click **Generate Self-signed Certificate Key** to have the WAAS devices use a self-signed certificate/key pair for SSL.
- Click **Import Existing Certificate Key** to upload or paste an existing certificate/key pair.
- Click **Export Certificate Key** to export the current certificate/key pair.
- Click **Generate Certificate Signing Request** to renew or replace the existing certificate/key pair. The certificate signing request (CSR) is used by the Certificate Authority to generate a new certificate.

  The file that you import or export must be in either a PKCS12 format or a PEM format.

For service certificate and private key configuration steps, see the "Configuring a Service Certificate and Private Key" section on page 13-32.

Step 8 Click **Submit**.

Configuring a Service Certificate and Private Key

To configure a service certificate and private key, follow these steps:

Step 1 To generate a self-signed certificate and private key (see Figure 13-17), follow these steps:

Figure 13-17 Self-Signed Certificate and Private Key

a. Check the **Mark private key as exportable** check box to export this certificate/key in the WAAS Central Manager and device CLI later.

b. Fill in the certificate and private key fields.

Step 2 To import an existing certificate or certificate chain and, optionally, private key (see Figure 13-18), follow these steps:

Note WAAS SSL feature only supports RSA signing/encryption algorithm and keys.

Figure 13-18 Importing Existing Certificate or Certificate Chain

a. Check the **Mark private key as exportable** check box to export this certificate/key in the WAAS Central Manager and device CLI later.

b. To import existing certificate or certificate chain and private key, perform one of the following:
   - Upload certificate and key in PKCS#12 format (also as Microsoft PFX format)
   - Upload certificate and private key in PEM format
   - Paste certificate and private key PEM content

   If the certificate and private key are already configured, you can update the certificate only. In this case, the Central Manager constructs the certificate and private key pair using the imported certificate and current private key. This functionality can be used to update an existing self-signed certificate to one signed by the Certificate Authority, or to update an expiring certificate.

   The Central Manager allows importing a certificate chain consisting of an end certificate that must be specified first, a chain of intermediate CA certificates that sign the end certificate or intermediate CA certificate, and end with a root CA.

   The Central Manager validates the chain and rejects it if the validity date of the CA certificate is expired, or the signing order of certificates in the chain is not consequent.

c. Enter a pass-phrase to decrypt the private key, or leave this field empty if the private key is not encrypted.

Step 3 To export a configured certificate and private key (see Figure 13-19), follow these steps:

Figure 13-19 Export Certificate and Key

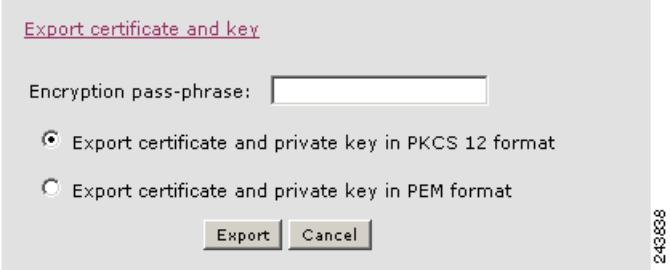

a. Enter the encryption pass-phrase.

b. Export current certificate and private key in either PKCS#12 or PEM formats. In case of PEM format both certificate and private key are included in single PEM file.

Note Central Manager will not allow exporting certificate and private key if the certificate and key were marked as non-exportable when they were generated or imported.

Step 4 To generate a certificate signing request from a current certificate and private key (see Figure 13-20), follow these steps:

Figure 13-20 Generate Certificate Signing Request

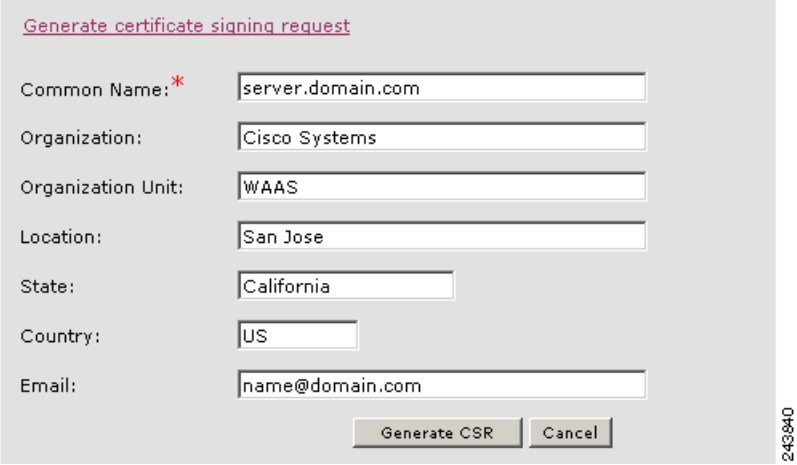

To update the current certificate with one signed by the Certificate Authority:

a. Generate PKCS#10 certificate signing request.

b. Send generated certificate signing request to Certificate Authority to generate and sign certificate.

c. Import certificate received from the Certificate Authority using the **Importing existing certificate and optionally private key** option.

Note The size of the key for a generated certificate request is the same as the size of the key in the current certificate.

Working with Cipher Lists

Cipher lists are sets of cipher suites that you can assign to your SSL acceleration configuration. A cipher suite is an SSL encryption method that includes the key exchange algorithm, the encryption algorithm, and the secure hash algorithm.

To configure a cipher list, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Security > SSL > Cipher Lists**.

The SSL Cipher Lists window appears (see Figure 13-21).

Note For a WAAS Express device, the SSL Cipher Lists window shows the same name and cipher fields but in a slightly different format.

Figure 13-21 SSL Cipher Lists Window

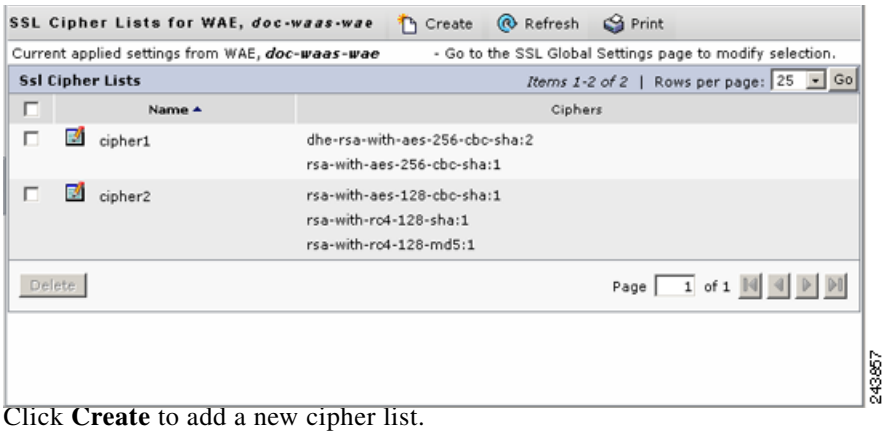

Step 3 Click **Create** to add a new cipher list.

The Creating New SSL Cipher List window appears (see Figure 13-22).

Note For a WAAS Express device, click **Add Cipher List** to add a new cipher list.

Figure 13-22 Creating New SSL Cipher List Window

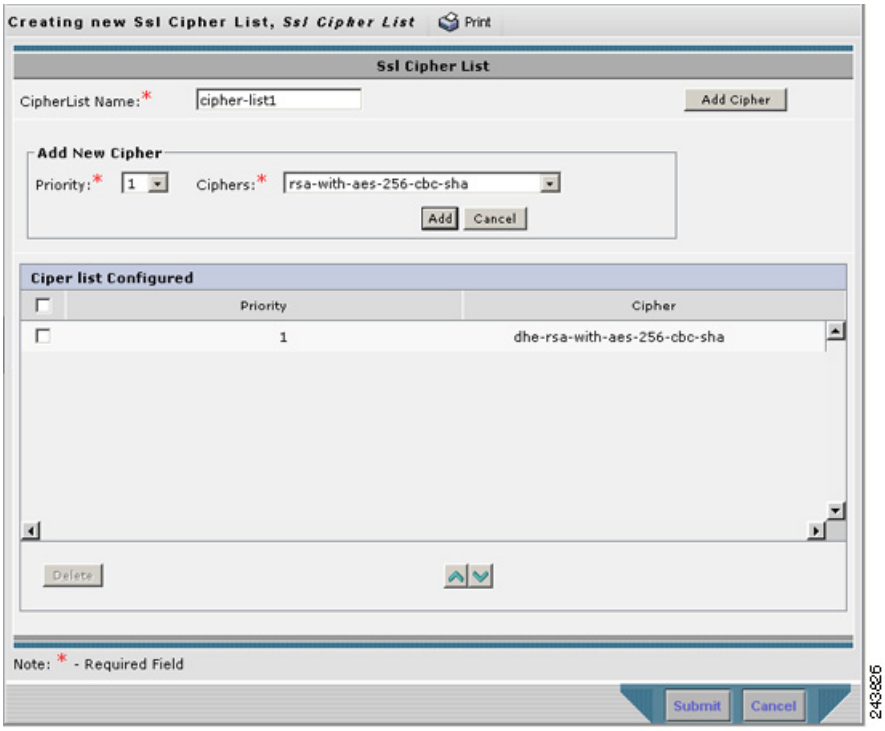

Step 4 Type a name for your cipher list in the Cipher List Name field.

Step 5 Click **Add Cipher** to add cipher suites to your cipher list.

Note For a WAAS Express device, select the ciphers you wish to add. Skip to Step 12.

Step 6 Choose the cipher suite that you want to add in the Ciphers field.

Note If you are establishing an SSL connection to a Microsoft IIS server, do not select a DHE-based cipher suite.

Step 7 Choose the priority for the selected cipher suite in the Priority field.

Note When SSL peering service is configured, the priority associated with a cipher list on a core device takes precedence over the priority associated with a cipher list on an edge device.

Step 8 Click **Add** to include the selected cipher suite on your cipher list, or click **Cancel** to leave the list as it is.

Step 9 Repeat Step 5 through Step 8 to add more cipher suites to your list as desired.

Step 10 (Optional) To change the priority of a cipher suite, check the cipher suite check box and then use the up or down arrow buttons located below the cipher list to prioritize.

Note The client-specified order for ciphers overrides the cipher list priority assigned here if the cipher list is applied to an accelerated service. The priorities assigned in this cipher list are only applicable if the cipher list is applied to SSL peering and management services.

Step 11 (Optional) To remove a cipher suite from the list, check the cipher suite's box and then click **Delete**.

Step 12 Click **Submit** when you are done configuring the cipher list.

Note For a WAAS Express device, click **OK** to save the cipher list configuration.

SSL configuration changes will not be applied on the device until the security license has been enabled on the device.

Working with Certificate Authorities

The WAAS SSL acceleration feature allows you to configure the Certificate Authority (CA) certificates used by your system. You can use one of the many well-known CA certificates that is included with WAAS or import your own CA certificate.

To manage your CA certificates, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Security > SSL > Certificate Authorities**.

The SSL CA Certificate List window appears (see Figure 13-23).

Note For a WAAS Express device, the SSL CA Certificate List window shows the same Name, Issued To, Issuer, Expiry Date fields but in a slightly different format.

There is also an **Aggregate Settings** field configurable as Yes or No. To finish the procedure for WAAS Express, skip to Step 4.

Figure 13-23 SSL CA Certificate List Window

Step 3 Add one of the preloaded CA certificates that is included with WAAS as follows:

a. Click **Well-known CAs**.

b. Choose the pre-existing CA certificate you want to add and click **Import**. The CA certificate that you selected is added to the list on the SSL CA Certificate List display.

Step 4 Add your own CA certificate as follows:

a. Click **Create**. The Creating New CA Certificate window appears (see Figure 13-24).

Note For a WAAS Express device, click **Add CA** to add your own CA certificate. Enter the name and the URL, then click **Get CA Certificate**. Skip to Step 6.

Figure 13-24 Creating New CA Certificate Window

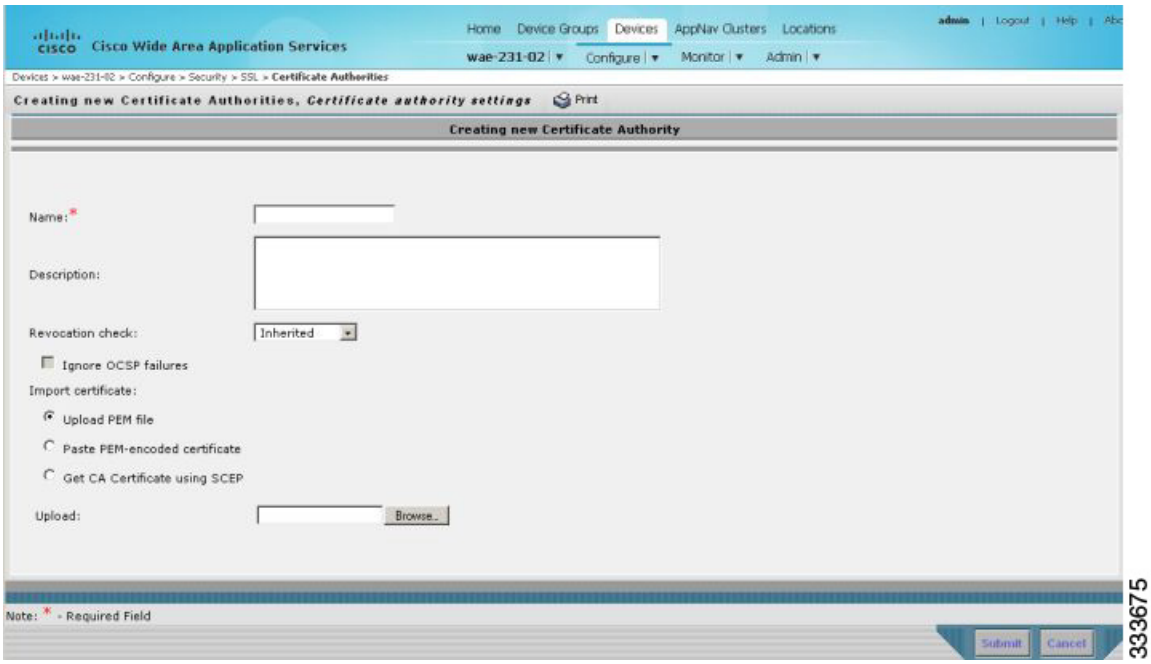

b. Type a name for the certificate in the Certificate Name field.

c. (Optional) Type a description of the CA certificate in the Description field.

d. Choose **disabled** in the Revocation check drop-down list to disable OCSP revocation of certificates signed by this CA. Check the **Ignore OCSP failures** check box to mark revocation check successful if the OCSP revocation check failed.

e. Add the certificate information by choosing on of the following methods:

   – **Upload PEM File**

     If you are uploading a file, it must be in a Privacy Enhanced Mail (PEM) format. Browse to the file that you want to use and click **Upload**.

   – **Paste PEM Encoded Certificate**

     If you are pasting the CA certificate information, paste the text of the PEM format certificate into the Paste PEM Encoded certificate field.

   – **Get CA Certificate using SCEP**

     This option automatically configures the certificate authority using Simple Certificate Enrollment Protocol. If you are using the automated certificate enrollment procedure, enter the CA URL and click **Get Certificate**. The contents of the certificate is displayed in text and PEM formats.

     To complete the automated certificate enrollment procedure, you must configure the SSL auto enrollment settings in the "SSL Auto Enrollment" section on page 13-40.

f. Click **Submit** to save your changes.

Step 5 (Optional) To remove a Certificate Authority from the list, select it and then click the **Delete** icon located in the toolbar.

Step 6 Click **Submit** when you are done configuring the CA certificate list.

Note For a WAAS Express device, click **OK** to save the CA certificate configuration.

SSL Auto Enrollment

The WAAS SSL acceleration feature allows you to enroll certificates automatically for a device (or device group) using SCEP. Once the CA certificate his been obtained, SSL auto enrollment settings must be configured.

Note You must configure the applicable certificate authority before configuring auto enrollment settings.

To configure SSL auto enrollment settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Security > SSL > Auto Enrollment**.

The SSL Auto Enrollment Settings window appears (see Figure 13-25).

Figure 13-25 SSL Auto Enrollment Settings Window

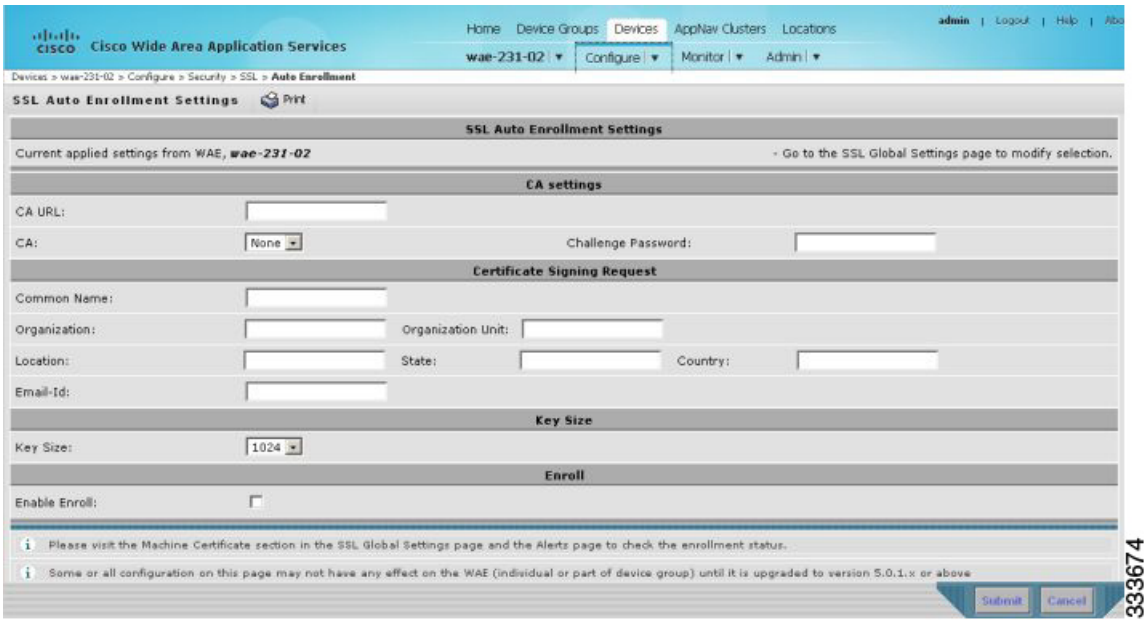

Step 3 Configure the following CA settings:

- CA URL
- CA—Select the appropriate CA from the list
- Challenge Password

Note CA, CA URL, and challenge password settings are mandatory for enabling SSL auto enrollment.

Step 4 Configure the following Certificate Signing Request settings:

- Common Name
- Organization and Organization Unit
- Location, State, and Country
- Email-Id

Step 5 Configure the key size: 512, 768, 1024, 1536, or 2048

Step 6 Check the Enable Enroll box.

Step 7 Click **Submit**.

You can then check the enrollment status in the Machine Certificate section on the SSL Global Settings page and on the Alerts page.

Configuring SSL Management Services

SSL management services are the SSL configuration parameters that affect secure communications between the Central Manager and the WAE devices (see Figure 13-14 on page 13-30). The certificate/key pairs used are unique for each WAAS device, and so SSL management services can only be configured for individual devices, not device groups.

To configure SSL management services, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Security > Management Service**.

The Management Services window appears (see Figure 13-26).

Figure 13-26 SSL Management Services Window

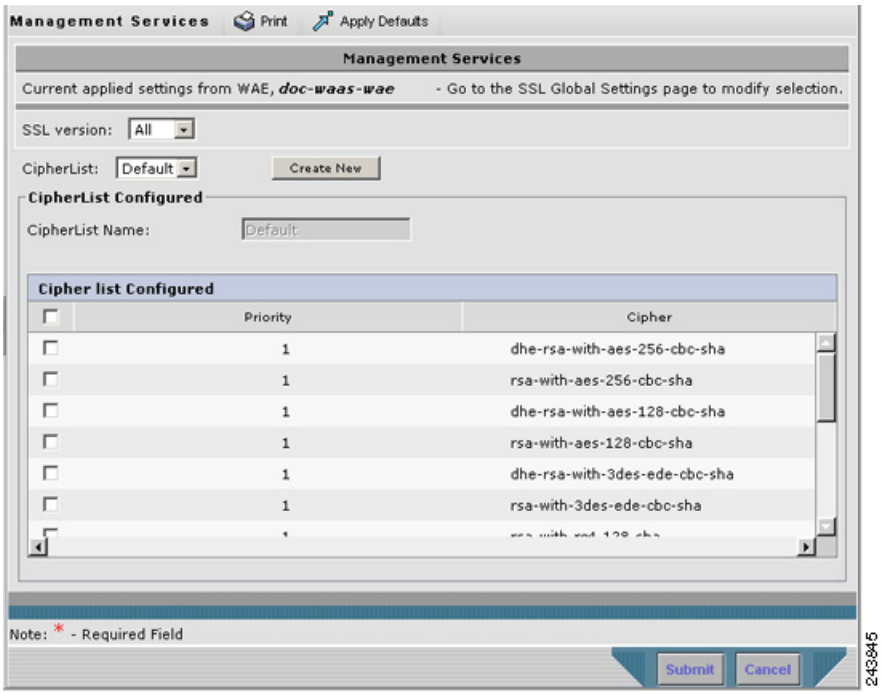

Step 3 In the SSL version field, choose the type of SSL protocol to use. Choose **SSL3** for the SSL version 3 protocol, choose **TLS1** for the Transport Layer Security version 1 protocol, or choose **All** to use both SSL3 and TLS1 SSL protocols.

Note Management service SSL version and cipher settings configured for the WAAS Central Manager are also applied to SSL connections between the WAAS Central Manager and the browser of the user.

Primary and standby Central Managers must share a common management service version or cipher list. Changing the management service version and cipher list settings may result in a loss of connectivity between primary Central Manager and standby Central Manager and WAE devices.

Table 13-3 shows the cipher lists supported with Internet Explorer and Mozilla Firefox:

Table 13-3 Cipher Lists Supported with Internet Explorer and Mozilla Firefox

Cipher	Internet Explorer	Firefox
dhe-rsa-with-aes-256-cbc-sha	Supported in IE7/Vista	Supported
rsa-with-aes-256-cbc-sha	Supported in IE7/Vista	Supported
dhe-rsa-with-aes-128-cbc-sha	Supported in IE7/Vista	Supported
rsa-with-aes-128-cbc-sha	Supported in IE7/Vista	Supported
dhe-rsa-with-3des-ede-cbc-sha	Not enabled by default	Supported
rsa-with-3des-ede-cbc-sha	Not enabled by default	Supported
rsa-with-rc4-128-sha	Supported	Supported

Table 13-3 Cipher Lists Supported with Internet Explorer and Mozilla Firefox

Cipher	Internet Explorer	Firefox
rsa-with-rc4-128-md5	Supported	Supported
dhe-rsa-with-des-cbc-sha	Not Supported	Not enabled by default
rsa-export1024-with-rc4-56-sha	Supported	Not enabled by default
rsa-export1024-with-des-cbc-sha	Supported	Not enabled by default
dhe-rsa-export-with-des40-cbc-sha	Not Supported	Not Supported
rsa-export-with-des40-cbc-sha	Not Supported	Not Supported
rsa-export-with-rc4-40-md5	Supported	Supported

Note Both Mozilla Firefox and Internet Explorer support SSLv3 and TLSv1 protocols, however TLSv1 may not be enabled by default. Therefore, you need to enable it in your browser.

Configuring ciphers or protocols that are not supported in your browser will result in connection loss between the browser and the Central Manager. If this occurs, configure the Central Manager management service SSL settings to the default in the CLI to restore the connection.

Some browsers, such as Internet Explorer, do not correctly handle a change of SSL version and cipher settings on the Central Manager, which can result in the browser showing an error page after submitting changes. If this occurs, reload the page.

Step 4 In the Cipher List pane, choose a list of cipher suites to be used for SSL acceleration. See the "Working with Cipher Lists" section on page 13-35 for additional information.

Configuring SSL Peering Service

SSL peering service configuration parameters control secure communications established by the SSL accelerator between WAE devices while optimizing SSL connections (see Figure 13-14 on page 13-30). The peering service certificate and private key is unique for each WAAS device and can only be configured for individual devices, not device groups.

To configure SSL peering service, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Configure > Security > Peering Service**.

The Peering Service window appears (see Figure 13-27).

Note For a WAAS Express device, the Peering Service window shows a subset of the fields in the standard Peering Service window in a slightly different format.

Cipher list **Priority** setting and **Disable revocation check of peer certificates** option are not applicable to WAAS Express.

Figure 13-27 SSL Peering Service Window

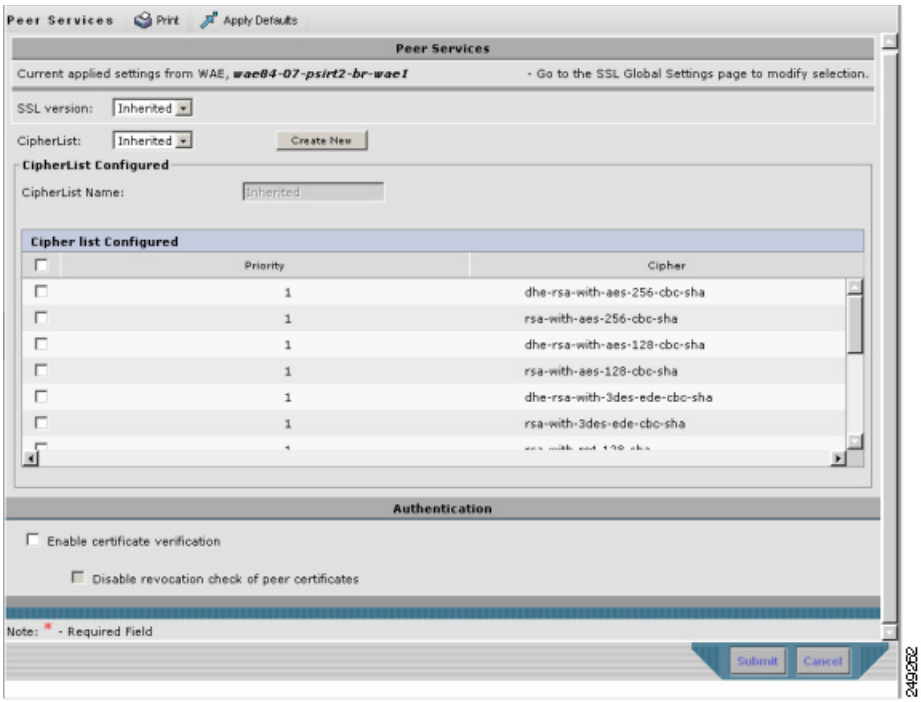

Step 3 In the SSL Version field, choose the type of SSL protocol to use, or choose **Inherited** to use the SSL protocol configured in global SSL settings. Choose **SSL3** for the SSL version 3 protocol, choose **TLS1** for the Transport Layer Security version 1 protocol, or choose **All** to use both SSL3 and TLS1 SSL protocols.

Note For a WAAS Express device, only SSL3 and TLS1 are supported for the SSL Version.

Step 4 To enable verification of peer certificates check **Enable Certificate Verification** check box. If certificate verification is enabled, WAAS devices that use self-signed certificates will not be able to establish peering connections to each other and, thus, not be able to accelerate SSL traffic.

Step 5 Check the **Disable revocation check for this service** check box to disable OCSP certificate revocation checking.

Note For a WAAS Express device, this option is not applicable.

Step 6 In the Cipher List pane, choose a list of cipher suites to be used for SSL acceleration between the WAE device peers, or choose **Inherited** to use the cipher list configured in SSL global settings.

Note For a WAAS Express device, the list of cipher suites to be used for SSL acceleration is shown in the Cipher List pane.

See the "Working with Cipher Lists" section on page 13-35 for additional information.

Step 7 Click **Submit**.

Note For a WAAS Express device, SSL configuration changes will not be applied on the device until the security license has been enabled on the device.

Using SSL Accelerated Services

After you have enabled and configured SSL acceleration on your WAAS system, you must define at least one service to be accelerated on the SSL path. To configure SSL accelerated services, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > SSL Accelerated Services**.

Step 3 To delete an accelerated service, select the service and click **Delete**.

Step 4 Click **Create** to define a new accelerated service. A maximum of 128 accelerated services are allowed. The Basic SSL Accelerated Services Configuration window appears (see Figure 13-28).

Figure 13-28 SSL Accelerated Services—Basic Window

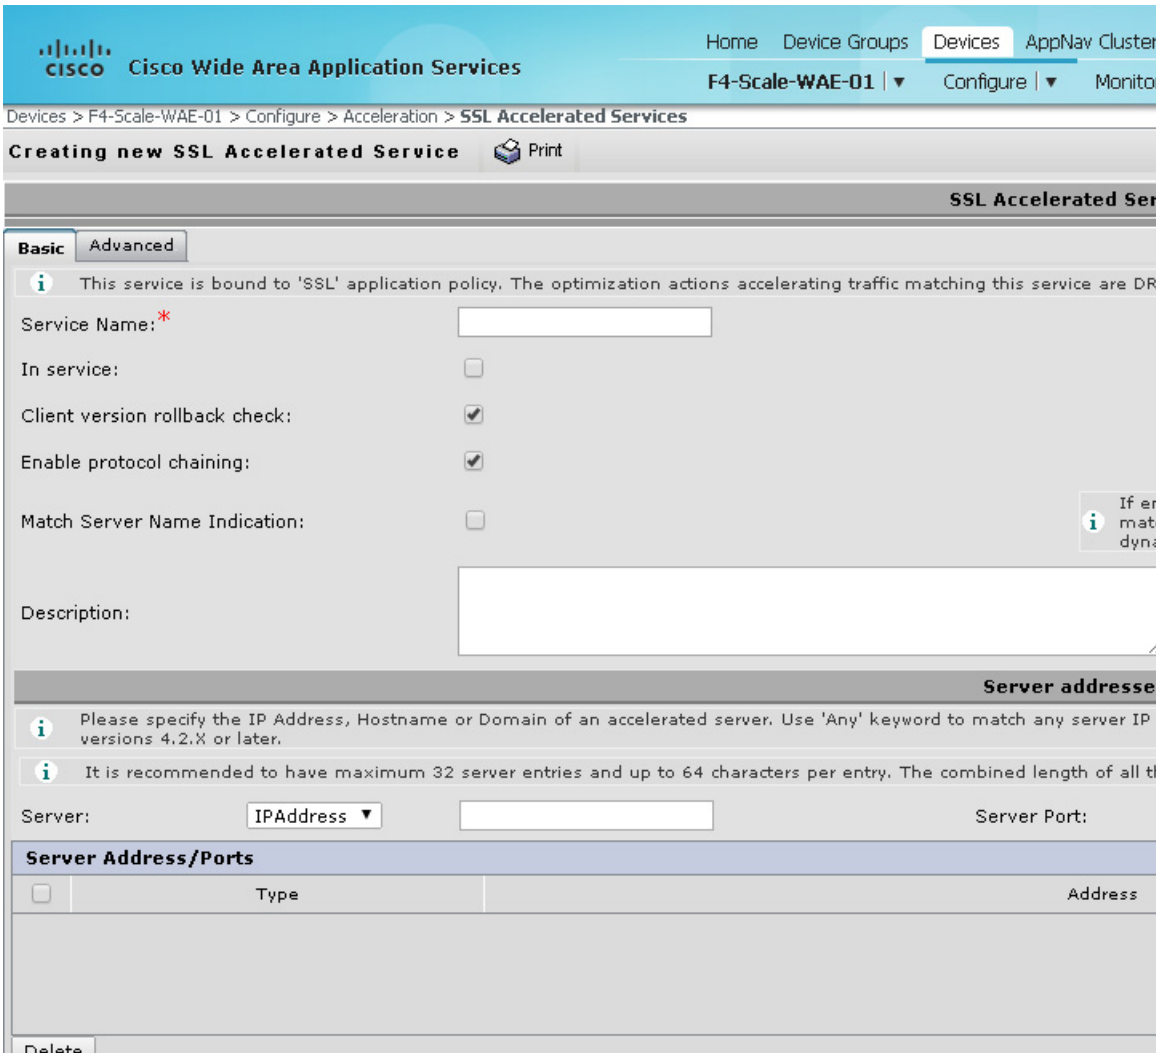

Step 5 Type a name for the service in the Service Name field.

Step 6 To enable this accelerated service, check the **In service** check box.

Step 7 To enable client version rollback check, check the **Client version rollback check** check box.

Enabling the client version rollback check does not allow connections with an incorrect client version to be optimized.

Step 8 (Optional) Type a description of the service in the Description field.

Step 9 From the Server drop-down list, choose **IP Address**, **Hostname**, or **Domain** as the SSL service endpoint type. Type the server IP address, hostname, or domain of the accelerated server. Use the keyword **Any** to specify any server IP address. A maximum of 32 IP addresses, 32 hostnames, and 32 domains are allowed.

Note Hostname and domain server address types are supported only when using WAAS software version 4.2.x or later. Server IP address keyword **Any** is supported only when using WAAS Software version 4.2.x or later.

Step 10 Type the port associated with the service to be accelerated. Click **Add** to add each address. If you specify a server hostname, the Central Manager resolves the hostname to the IP address and adds it to the Server IP/Ports table.

Step 11 Click **Delete** to remove an IP address from the list.

Step 12 Choose a certificate and key pair method (see Figure 13-29).

Figure 13-29 Configuring Service Certificate and Private Key

- Click **Generate Self-signed Certificate Key** to have the WAAS devices use a self-signed certificate/key pair for SSL.
- Click **Import Existing Certificate Key** to upload or paste an existing certificate/key pair.
- Click **Export Certificate Key** to export the current certificate/key pair.
- Click **Generate Certificate Signing Request** to renew or replace the existing certificate/key pair. The certificate signing request (CSR) is used by the Certificate Authority to generate a new certificate.

  The file that you import or export must be in either a PKCS12 format or a PEM format.

For service certificate and private key configuration steps, see the "Configuring a Service Certificate and Private Key" section on page 13-32.

Note If you change the certificate or key for an existing SSL accelerated service, you must uncheck the **In service** check box and click **Submit** to disable the service, then wait 5 minutes and check the **In service** check box and click **Submit** to reenable the service. Alternatively, at the WAE, you can use the **no inservice** SSL accelerated service configuration command, wait a few seconds, and then use the **inservice** command. If you are changing the certificate or key for multiple SSL accelerated services, you can restart all accelerated services by disabling and then reenabling the SSL accelerator.

Step 13 Click the **Advanced Settings** tab to configure SSL parameters for the service. The Advanced SSL Accelerated Services Configuration window appears (see Figure 13-30).

Figure 13-30 SSL Accelerated Services—Advanced Window

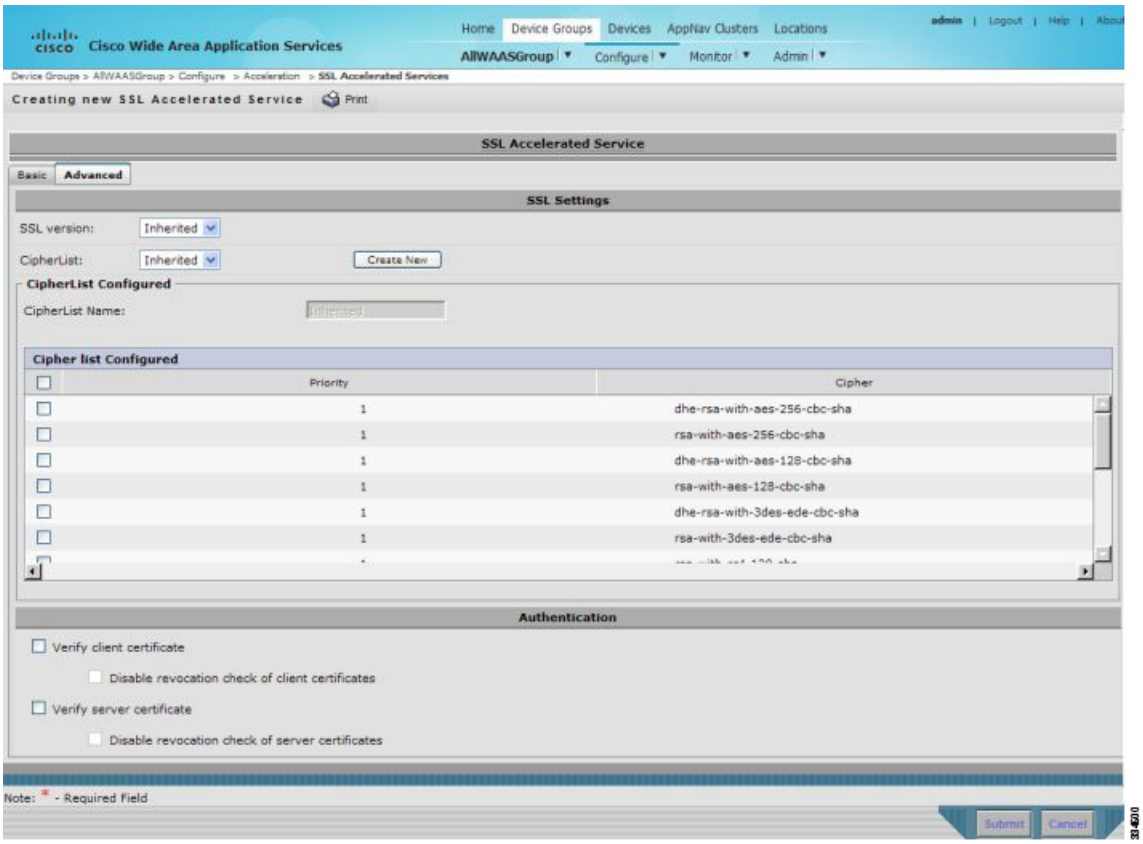

Step 14 (Optional) In the SSL version field, choose the type of SSL protocol to use, or choose **Inherited** to use the SSL protocol configured in global SSL settings. Choose **SSL3** for the SSL version 3 protocol, choose **TLS1** for the Transport Layer Security version 1 protocol, or choose **All** to use both SSL3 and TLS1 SSL protocols.

Step 15 (Optional) In the Cipher List field, choose a list of cipher suites to be used for SSL acceleration between the WAE device peers, or choose **Inherited** to use the cipher list configured in SSL global settings. For more information, see the "Working with Cipher Lists" section on page 13-35.

Step 16 (Optional) To set the Online Certificate Status Protocol (OCSP) parameters for certificate revocation, follow these steps:

a. To enable verification of client certificate check, check the **Verify client certificate** check box.

b. Check the **Disable revocation check for this service** check box to disable OCSP client certificate revocation checking.

c. To enable verification of server certificate check, check the **Verify server certificate** check box.

d. Check the **Disable revocation check for this service** check box to disable OCSP server certificate revocation checking.

Note If the server and client devices are using self-signed certificates and certificate verification is enabled, WAAS devices will not be able to accelerate SSL traffic.

Step 17 Click **Submit** when you have finished configuring the SSL accelerated service.

Creating a New Traffic Optimization Policy

Table 13-4 provides an overview of the steps that you must complete to create a new traffic optimization policy.

Table 13-4 Checklist for Creating a New Optimization Policy

Task	Additional Information and Instructions
1. Prepare for creating an optimization policy.	Provides the tasks you need to complete before creating a new optimization policy on your WAAS devices. For more information, see the "Preparing to Create an Optimization Policy" section on page 13-49.
2. Create an application definition.	Identifies general information about the application you want to optimize, such as the application name and whether the WAAS Central Manager collects statistics about this application. For more information, see the "Creating an Application Definition" section on page 13-50.
3. Create an optimization policy.	Determines the type of action your WAAS device or device group performs on specific application traffic. This step requires you to do the following: • Create application class maps that allow a WAAS device to identify specific types of traffic. For example, you can create a condition that matches all traffic going to a specific IP address. • Specify the type of action your WAAS device or device group performs on the defined traffic. For example, you can specify that WAAS should apply TFO and LZ compression to all traffic for a specific application. For more information, see the "Creating an Optimization Policy" section on page 13-51.

Preparing to Create an Optimization Policy

Before you create a new optimization policy, complete the following preparation tasks:

- Review the list of optimization policies on your WAAS system and make sure that none of these policies already cover the type of traffic you want to define. To view a list of the predefined policies that come bundled with the WAAS system, see Appendix A, "Predefined Optimization Policy."
- Identify a match condition for the new application traffic. For example, if the application uses a specific destination or source port, you can use that port number to create a match condition. You can also use a source or destination IP address for a match condition.
- Identify the device or device group that requires the new optimization policy. We recommend you create optimization policies on device groups so the policy is consistent across multiple WAAS devices.

Creating an Application Definition

The first step in creating an optimization policy is to set up an application definition that identifies general information about the application, such as the application name and whether you want the WAAS Central Manager to collect statistics about the application. You can create up to 255 application definitions on your WAAS system.

To create an application definition, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Configure > Acceleration > Applications**.

The Applications window appears, which displays a list of all applications on your WAAS system. It also lists the device or device group from which it gets the settings. From this window, you can perform the following tasks:

- Select an application and click the **Edit** icon in the task bar to modify or click the **Delete** icon in the task bar to delete.
- Determine if your WAAS system is collecting statistics on an application. The Enable Statistics column displays Yes if statistics are being collected for the application.
- Create a new application as described in the steps that follow.

Click the **Add Application** icon in the taskbar. The Application window appears.

Step 2 Enter a name for this application.

The name cannot contain spaces and special characters.

Step 3 (Optional) Enter a comment in the **Comments** field.

The comment you enter appears in the Applications window.

Step 4 Check the **Enable Statistics** check box to allow the WAAS Central Manager to collect data for this application. To disable data collection for this application, uncheck this box.

The WAAS Central Manager GUI can display statistics for up to 25 applications and 25 class maps. An error message is displayed if you try to enable more than 25 statistics for either. However, you can use the WAAS CLI to view statistics for all applications that have policies on a specific WAAS device. For more information, refer to the *Cisco Wide Area Application Services Command Reference*.

If you are collecting statistics for an application and decide to disable statistics collection, then reenable statistics collection at a later time, the historical data will be retained, but a gap in data will exist for the time period when statistics collection was disabled. An application cannot be deleted if there is an optimization policy using it. However, if you delete an application that you had collected statistics for, then later recreate the application, the historical data for the application will be lost. Only data since the recreation of the application will be displayed.

Note The WAAS Central Manager does not start collecting data for this application until you finish creating the entire optimization policy.

Step 5 Click **OK**.

The application definition is saved and is displayed in the application list.

Creating an Optimization Policy

After you create an application definition, you need to create an optimization policy that determines the action a WAAS device takes on the specified traffic. For example, you can create an optimization policy that makes a WAAS device apply TCP optimization and compression to all application traffic that travels over a specific port or to a specific IP address. You can create up to 512 optimization policies on your WAAS system.

The traffic matching rules are contained in the application class map. These rules, known as match conditions, use Layer 2 and Layer 4 information in the TCP header to identify traffic.

To create an optimization policy, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > Optimization Policies**.

The Optimization Policies window appears. (See Figure 13-31.)

Note For a WAAS Express device, the Optimization Policies window shows a subset of the fields in the standard Optimization Policies window.

Enable Service Policy option, **DSCP** option, and the **Protocol** column in the list of policy rules are not applicable to WAAS Express.

Figure 13-31 Optimization Policies Window

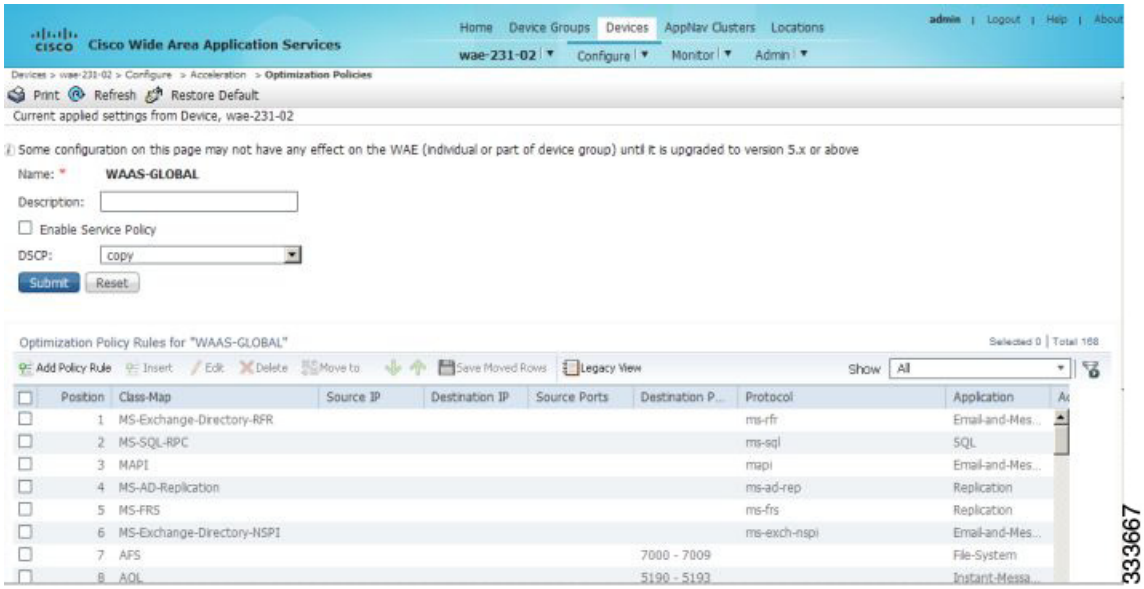

This window displays information about all optimization policies that reside on the selected device or device group and the position of each policy. The position determines the order in which WAAS refers to that policy when determining how to handle application traffic. To change the position of a policy, see the "Modifying the Position of an Optimization Policy" section on page 13-59. This window also displays the class map, source and destination IP addresses, source and destination ports, protocol, application, action, and accelerate assigned to each policy.

Note If there are version 4.x devices, you can click the **Legacy View** taskbar icon to view the policies as they appear in a 4.x device.

From the Optimization Policies window, you can perform the following tasks:

- Configure a description, configure the Enable Service Policy setting, and configure the DSCP setting. This DSCP setting field configures DSCP settings at the device (or device group) level.

  Note The device will only use this policy setting to determine what optimizations are done if Enable Service Policy is set.

- Select one or more optimization policies that you want to delete, and click the **Delete** icon to delete the checked policies.
- Select an optimization policy and click the **Edit** icon to modify the checked policy.
- Restore predefined policies and class maps. For more information, see the "Restoring Optimization Policies and Class Maps" section on page 13-58.
- Create an optimization policy as described in the steps that follow.

Step 3 Click the **Add Policy Rule** icon in the taskbar to create a new optimization policy.

The Optimization Policy Rule pop-up window appears. (See Figure 13-32.)

Figure 13-32 Add Optimization Policy Rule Window

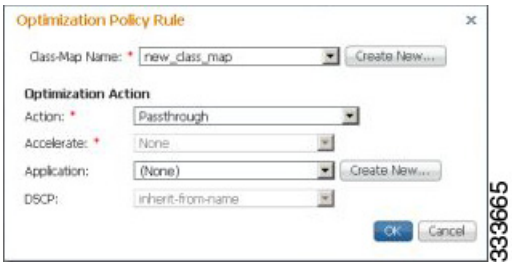

Step 4 Choose the class map from the Class-Map Name drop-down list to select an existing class map for this policy or click **Create New** to create a new class map for this policy.

If you are selecting an existing class map, skip the next step.

Step 5 Complete the following steps to create a new class map:

a. Enter a name for this application class map. The name cannot contain spaces or special characters.

Note For WAAS Express, the class map name cannot contain the following prefixes (case sensitive): Class, optimize, passthrough, application, accelerate, tfo, dre, lz, or sequence-interval. Existing class map names containing any of these prefixes must be changed manually.

b. (Optional) Enter a description.

c. From the Type drop-down list, choose the class map type. Once you have chosen the type, you can enter the match conditions.

d. Click the **Add Match Condition** icon to enter the conditions. (See Figure 13-33.)

Note For a WAAS Express device, **Protocol** and **EPM Custom UUID** settings are not applicable.

Figure 13-33 Adding a New Match Condition Window

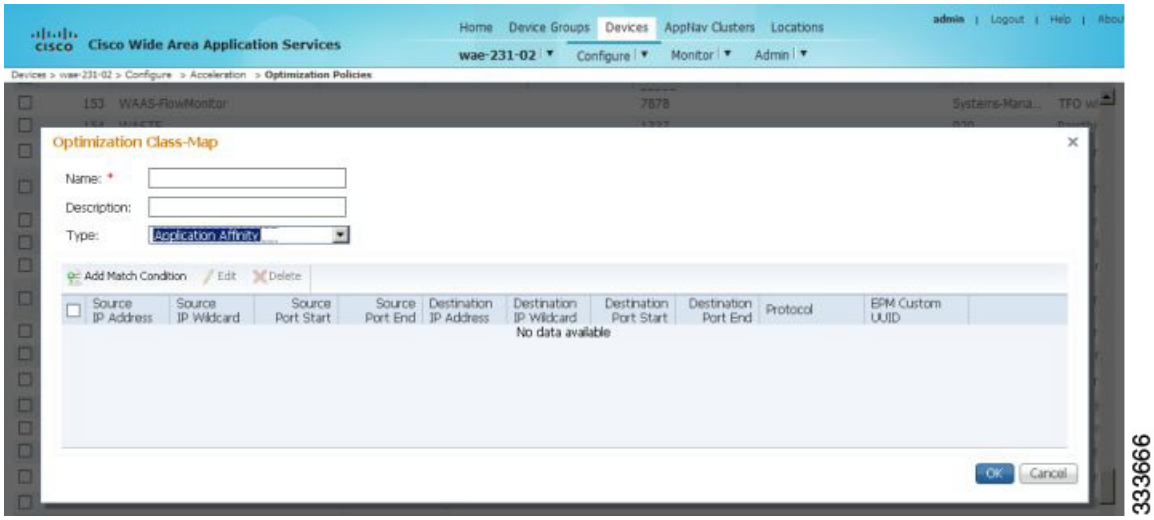

e. Enter a value in one of the destination or source condition fields to create a condition for a specific type of traffic.

For example, to match all traffic going to IP address 10.10.10.2, enter that IP address in the Destination IP Address field.

Note To specify a range of IP addresses, enter a wildcard subnet mask in either the destination or source IP Wildcard field.

f. Click **OK**. You return to creating a new policy in the Optimization Policy window.

Step 6 From the Action drop-down list, choose the action that your WAAS device should take on the defined traffic. Table 13-5 describes each action.

Note For a WAAS Express device, only a subset of the actions are available. These include: Passthrough, TFO Only, TFO with LZ, TFO with DRE, and TFO with DRE and LZ.

Table 13-5 Action Descriptions

Action[1]	Description
Passthrough	Prevents the WAAS device from optimizing the application traffic defined in this policy by using TFO, DRE, or compression. Traffic that matches this policy can still be accelerated if an accelerator is chosen from the Accelerate drop-down list.
TFO Only	Applies a variety of transport flow optimization (TFO) techniques to matching traffic. TFO techniques include BIC-TCP, window size maximization and scaling, and selective acknowledgement. For a more detailed description of the TFO features, see the "TFO Optimization" section on page 1-4.
TFO with DRE (Adaptive Cache)	Applies both TFO and DRE with adaptive caching to matching traffic.
TFO with DRE (Unidirectional Cache)	Applies both TFO and DRE with unidirectional caching to matching traffic.
TFO with DRE (Bidirectional Cache)	Applies both TFO and DRE with bidirectional caching to matching traffic.
TFO with LZ Compression	Applies both TFO and the LZ compression algorithm to matching traffic. LZ compression functions similarly to DRE but uses a different compression algorithm to compress smaller data streams and maintains a limited compression history.
TFO with DRE (Adaptive Cache) and LZ	Applies TFO, DRE with adaptive caching, and LZ compression to matching traffic.
TFO with DRE (Unidirectional Cache) and LZ	Applies TFO, DRE with unidirectional caching, and LZ compression to matching traffic.
TFO with DRE (Bidirectional Cache) and LZ	Applies TFO, DRE with bidirectional caching, and LZ compression to matching traffic.

1. When configuring a device running a WAAS version prior to 4.4.1, options that include Unidirectional or Adaptive caching are not shown in the Action list.

Note When configuring optimization policies on a device group, if the device group contains devices running a WAAS version prior to 4.4.1 and you are configuring an action that includes Unidirectional or Adaptive caching, the caching mode is converted to bidirectional. Similarly, when devices running WAAS versions prior to 4.4.1 join a device group that is configured with optimization policies that use Unidirectional or Adaptive caching, the caching mode is converted to bidirectional. In such cases, we recommend that you upgrade all devices to the same software version or create different device groups for devices with incompatible versions.

Step 7 From the Accelerate drop-down list, choose one of the following additional acceleration actions that your WAAS device should take on the defined traffic:

- **None**—No additional acceleration is done.
- **MS PortMapper**—Accelerate using the Microsoft Endpoint Port Mapper (EPM).
- **CIFS Adaptor**—Accelerate using the CIFS or SMB Accelerators.

- **HTTP Adaptor**—Accelerate using the HTTP Accelerator.
- **MAPI Adaptor**—Accelerate using the MAPI Accelerator.
- **NFS Adaptor**—Accelerate using the NFS Accelerator.
- **Video Adaptor**—Accelerate using the Video Accelerator.
- **ICA Adaptor**—Accelerate using the ICA Accelerator.

Note For a WAAS Express device, the available accelerators are CIFS Express and HTTP Express.

Step 8 Specify the application that you want to be associated with this policy by doing either of the following:

- From the Application drop-down list, choose an existing application like the one that you created in the "Creating an Application Definition" section on page 13-50. This list displays all predefined and new applications on your WAAS system.
- Click **New Application** to create an application. You can specify the application name and enable statistics collection. After specifying the application details, click **OK** to save the new application and return to the Optimization Policy window. The new application is automatically assigned to this device or device group.

Step 9 (Optional) Choose a value from the DSCP Marking drop-down list. You can choose copy, which copies the DSCP value from the incoming packet and uses it for the outgoing packet. If you choose inherit-from-name from the drop-down list, the DSCP value defined at the application or global level is used.

DSCP is a field in an IP packet that enables different levels of service to be assigned to network traffic. Levels of service are assigned by marking each packet on the network with a DSCP code and associating a corresponding level of service. DSCP is the combination of IP Precedence and Type of Service (ToS) fields. For more information, see RFC 2474.

DSCP marking does not apply to pass-through traffic.

Note For a WAAS Express device, the DSCP Marking drop-down list is not shown.

For the DSCP marking value, you can choose to use the global default values (see the "Defining Default DSCP Marking Values" section on page 13-58) or select one of the other defined values. You can choose copy, which copies the DSCP value from the incoming packet and uses it for the outgoing packet.

Step 10 Click **OK**.

The new policy appears in the Optimization Policies window. (See Figure 13-31 on page 13-51.)

Managing Application Acceleration

This section contains the following topics:

Modifying the Accelerator Load Indicator Threshold

To modify the accelerator load indicator threshold for a WAE device or device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > Accelerator Threshold**. The Accelerator Threshold window appears.

Step 3 In the **Accelerator Load Indicator Threshold** field, enter a percent value between 80 and 100. The default is 95.

Step 4 Click **Submit**.

Viewing a List of Applications

To view a list of applications that reside on a WAE device or device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > Optimization Policies**. The Optimization Policies window appears.

Step 3 Click the Application column header to sort the column by application name so you can more easily locate a specific application.

Note If there are version 4.x devices, you can click the **Legacy View** taskbar icon to view the policies as they appear in a 4.x device.

To edit an optimization policy, check the box next to the application and click the **Edit** taskbar icon.

If you determine that one or more policies are not needed, check the box next to each unneeded application and click the **Delete** taskbar icon.

If you determine that a new policy is needed, click the **Add Policy Rule** taskbar icon to create the policy (see the "Creating an Optimization Policy" section on page 13-51).

Viewing a Policy Report

To view a report of the policies that reside on each WAE device or device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Configure > Acceleration > Optimization Policy Report**. (See Figure 13-34.)

The Policy Report for Devices tab appears. This report lists each device (or device group) and the overall policy count on the device (or device group) referencing this application. It includes both active policies (those in use by the device or device group), and backup policies (those not in use by the device when the device gets its config from a device group). When the device is deassigned from the device group, the backup policies are applied back to the device and become active again.

An application cannot be deleted unless the No of Policies field is 0.

Figure 13-34 Optimization Policy Report

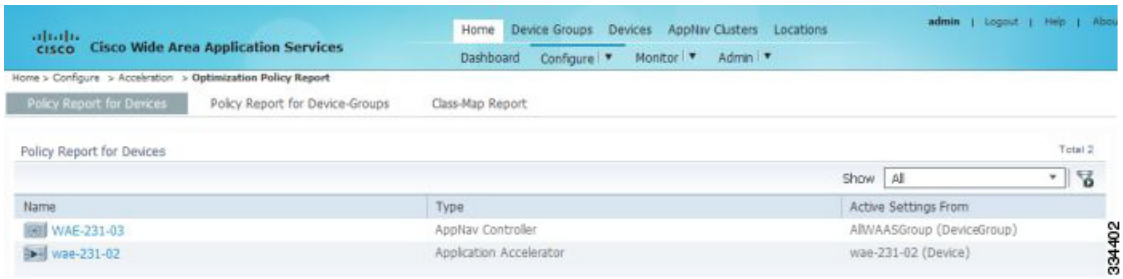

Step 2 Select the **Policy Report for Device-Groups** tab to view the number of devices per device group and the number of active policies in the device group.

Step 3 To see the optimization policies that are defined on a particular device or group, click the device or group to view the policies in the Optimization Policies window.

To view a class map report, see the "Viewing a Class Map Report" section on page 13-57.

Viewing a Class Map Report

To view a report of the class maps that reside on each WAE device or device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Configure > Acceleration > Optimization Policy Report**.

The Policy Report for Devices tab appears.

Step 2 Click the **Class-Map Report** tab to view a report of the devices and device groups on which the class map is configured.

Step 3 Select the class map and click the **View** icon to see the devices or device groups on which the class maps reside.

Restoring Optimization Policies and Class Maps

The WAAS system allows you to restore the predefined policies and class maps that shipped with the WAAS system. For a list of the predefined policies, see Appendix A, "Predefined Optimization Policy."

If you made changes to the predefined policies that have negatively impacted how a WAAS device handles application traffic, you can override your changes by restoring the predefined policy settings.

To restore predefined policies and class maps, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > Optimization Policies**.

The Optimization Policies window appears.

Step 3 Click the **Restore Default** taskbar icon to restore over 150 policies and class maps that shipped with the WAAS software and remove any new policies that were created on the system. If a predefined policy has been changed, these changes are lost and the original settings are restored.

Monitoring Applications and Class Maps

After you create an optimization policy, you should monitor the associated application to make sure your WAAS system is handling the application traffic as expected.

To monitor an application, you must have enabled statistics collection for that application, as described in the "Creating an Application Definition" section on page 13-50.

To monitor a class map, from the WAAS Central Manager menu, choose **Configure > Acceleration > Monitor Classmaps**. Select the class map on which to enable statistics and click the **Enable** button.

The WAAS Central Manager GUI can display statistics for up to 25 applications and 25 class maps. An error message is displayed if you try to enable more than 25 statistics for either. However, you can use the WAAS CLI to view statistics for all applications that have policies on a specific WAAS device. For more information, refer to the *Cisco Wide Area Application Services Command Reference*.

You can use the TCP Summary report to monitor a specific application. For more information, see the "TCP Summary Report" section on page 17-36.

Most charts can be configured to display Class Map data by clicking the chart Edit icon and choosing the Classifier series.

Defining Default DSCP Marking Values

According to policies that you define in an application definition and an optimization policy, the WAAS software allows you to set a DSCP value on packets that it processes.

A DSCP value is a field in an IP packet that enables different levels of service to be assigned to the network traffic. The levels of service are assigned by marking each packet on the network with a DSCP code and associating a corresponding level of service. The DSCP marking determines how packets for a connection are processed externally to WAAS. DSCP is the combination of IP Precedence and Type of Service (ToS) fields. For more information, see RFC 2474. DSCP values are predefined and cannot be changed.

This attribute can be defined at the following levels:

- Global—You can define global defaults for the DSCP value for each device (or device group) in the Optimization Policies page for that device (or device group). This value applies to the traffic if a lower level value is not defined.
- Policy—You can define the DSCP value in an optimization policy. This value applies only to traffic that matches the class maps defined in the policy and overrides the application or global DSCP value.

This section contains the following topic:

- Defining the Default DSCP Marking Value, page 13-59

Defining the Default DSCP Marking Value

To define the global default DSCP marking value, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > Optimization Policies**. The Optimization Policies window appears.

Step 3 Choose a value from the DSCP drop-down list. The default setting is copy, which copies the DSCP value from the incoming packet and uses it for the outgoing packet.

Step 4 Click **OK** to save the settings.

Modifying the Position of an Optimization Policy

Each optimization policy has an assigned position that determines the order in which a WAAS device refers to the policy in an attempt to classify traffic. For example, when a WAAS device intercepts traffic, it refers to the first policy in the list to try to match the traffic to an application. If the first policy does not provide a match, the WAAS device moves on to the next policy in the list.

You should consider the position of policies that pass through traffic unoptimized because placing these policies at the top of the list can cancel out optimization policies that appear farther down the list. For example, if you have two optimization policies that match traffic going to IP address 10.10.10.2, and one policy optimizes this traffic and a second policy in a higher position passes through this traffic, then all traffic going to 10.10.10.2 will go through the WAAS system unoptimized. For this reason, you should make sure that your policies do not have overlapping matching conditions, and you should monitor the applications you create to make sure that WAAS is handling the traffic as expected. For more information on monitoring applications, see Chapter 17, "Monitoring and Troubleshooting Your WAAS Network."

To modify the position of an optimization policy, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > Optimization Policies**. The Optimization Policies window appears. (See Figure 13-35.)

Note For a WAAS Express device, all policies are grouped under the waas_global category.

For a list of predefined policies, see Appendix A, "Predefined Optimization Policy."

Figure 13-35 Optimization Policies Window

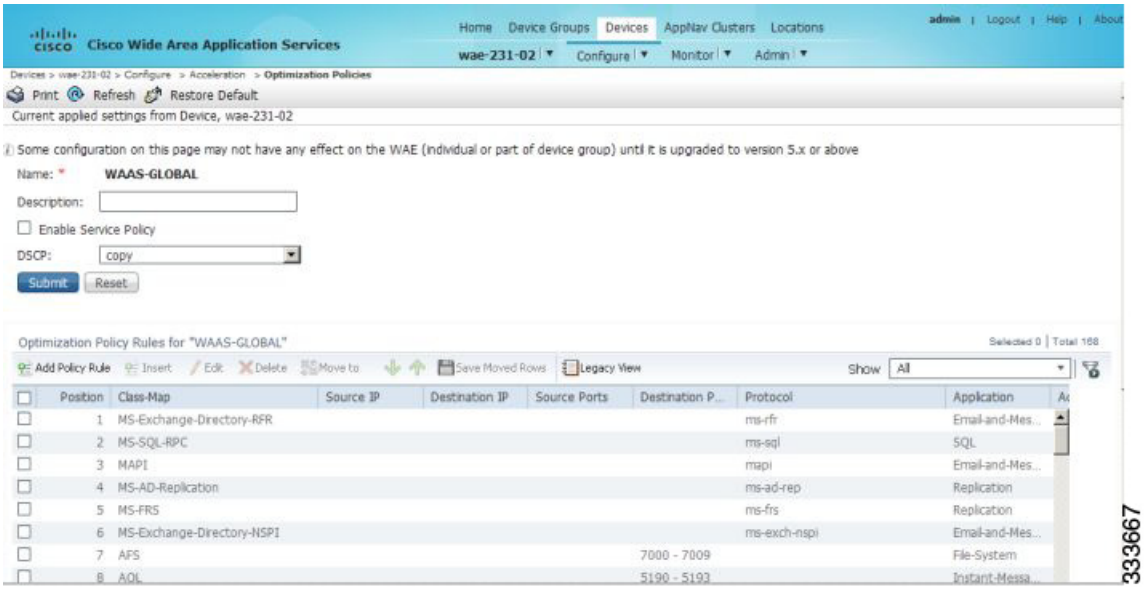

Step 3 Modify the position of the optimization policy in any of the following ways:

- Select the policy you would like to move and use the up and down arrow ( ) icons in the taskbar to move that policy higher or lower in the list.
- Select the policy you would like to move and use the Move To button to specify the exact position.
- Select the policy and drag and drop it into the desired position

Note The **Save Moved Rows** icon must be clicked to save the new policy positions.

You can also create a new optimization policy at a particular position by selecting the policy above the location and then clicking the Insert button.

If a device goes through all the policies in the list without making a match, then the WAAS device passes through the traffic unoptimized.

Note For a WAAS Express device, the class-default policy should be last. This policy cannot be modified or deleted.

Step 4 Click the **Save Moved Rows** icon to save any changes you made to policy positions.

Step 5 If you determine that a policy is not needed, follow these steps to delete the policy:

a. Select the policy you want to delete.

b. Click the **Delete** icon in the taskbar.

Note A default policy which maps to a default class map matching any traffic cannot be deleted.

Step 6 If you determine that a new policy is needed, click the **Add Policy** taskbar icon to create the policy (see the "Creating an Optimization Policy" section on page 13-51).

Modifying the Acceleration TCP Settings

In most cases, you do not need to modify the acceleration TCP settings because your WAAS system automatically configures the acceleration TCP settings based on the hardware platform of the WAE device. WAAS automatically configures the settings only under the following circumstances:

- When you first install the WAE device in your network.
- When you enter the **restore factory-default** command on the device. For more information about this command, see the *Cisco Wide Area Application Services Command Reference*.

The WAAS system automatically adjusts the maximum segment size (MSS) to match the advertised MSS of the client or server for each connection. The WAAS system uses the lower of 1432 or the MSS value advertised by the client or server.

If your network has high BDP links, you may need to adjust the default buffer settings automatically configured for your WAE device. For more information, see the "Calculating the TCP Buffers for High BDP Links" section on page 13-62.

If you want to adjust the default TCP adaptive buffering settings for your WAE device, see the "Modifying the TCP Adaptive Buffering Settings" section on page 13-63.

To modify the acceleration TCP settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > TCP Settings**. The Acceleration TCP Settings window appears.

Step 3 Keep the **Send TCP Keepalive** check box checked.

Checking the Send TCP Keepalive check box allows this WAE device or group to disconnect the TCP connection to its peer device if no response is received from the TCP keepalive exchange. In this case, the two peer WAE devices will exchange TCP keepalives on a TCP connection and if no response is received for the keepalives for a specific period, the TCP connection will be torn down. When the keepalive option is enabled, any short network disruption in the WAN will cause the TCP connection between peer WAE devices to be disconnected.

If the Send TCP Keepalive check box is not checked, TCP keepalives will not be sent and connections will be maintained unless they are explicitly disconnected. By default, this setting is enabled.

Step 4 Modify the TCP acceleration settings as needed. See Table 13-6 for a description of these settings.

For information on how to calculate these settings for high BDP links, see the "Calculating the TCP Buffers for High BDP Links" section on page 13-62.

Table 13-6 TCP Settings

TCP Setting	Description
Optimized Side	
Maximum Segment Size	Maximum packet size allowed between this WAAS device and other WAAS devices participating in the optimized connection. The default is 1432 bytes.
Send Buffer Size	Allowed TCP sending buffer size (in kilobytes) for TCP packets sent from this WAAS device to other WAAS devices participating in the optimized connection. The default is 32 KB.
Receive Buffer Size	Allowed TCP receiving buffer size (in kilobytes) for incoming TCP packets from other WAAS devices participating in the optimized connection. The default is 32 KB.
Original Side	
Maximum Segment Size	Maximum packet size allowed between the origin client or server and this WAAS device. The default is 1432 bytes.
Send Buffer Size	Allowed TCP sending buffers size (in kilobytes) for TCP packets sent from this WAAS device to the origin client or server. The default is 32 KB.
Receive Buffer Size	Allowed TCP receiving buffer size (in kilobytes) for incoming TCP packets from the origin client or server. The default is 32 KB.

Step 5 If you are deploying the WAE across a high Bandwidth-Delay-Product (BDP) link, you can set recommended values for the send and receive buffer sizes by clicking the **Set High BDP recommended values** button. For more information about calculating TCP buffers for high BDP links, see the "Calculating the TCP Buffers for High BDP Links" section on page 13-62.

Step 6 Click **Submit**.

Note If the original and optimized maximum segment sizes are set to their default values and you configure a jumbo MTU setting, the segment sizes are changed to the jumbo MTU setting minus 68 bytes. If you have configured custom maximum segment sizes, their values are not changed if you configure a jumbo MTU. For more information on jumbo MTU, see the "Configuring a Jumbo MTU" section on page 6-21.

To configure TCP keepalives from the CLI, use the **tfo tcp keepalive** global configuration command.

To configure TCP acceleration settings from the CLI, use the following global configuration commands: **tfo tcp optimized-mss**, **tfo tcp optimized-receive-buffer**, **tfo tcp optimized-send-buffer**, **tfo tcp original-mss**, **tfo tcp original-receive-buffer**, and **tfo tcp original-send-buffer**.

To show the TCP buffer sizes, use the **show tfo tcp** EXEC command.

Calculating the TCP Buffers for High BDP Links

WAAS software can be deployed in different network environments, involving multiple link characteristics such as bandwidth, latency, and packet loss. All WAAS devices are configured to accommodate networks with maximum Bandwidth-Delay-Product (BDP) of up to the values listed below:

- WAE-512—Default BDP is 32 KB
- WAE-612—Default BDP is 512 KB
- WAE-674 —Default BDP is 2048 KB
- WAE-7341 —Default BDP is 2048 KB
- WAE-7371 —Default BDP is 2048 KB
- All WAVE platforms—Default BDP is 2048 KB

If your network provides higher bandwidth or higher latencies are involved, use the following formula to calculate the actual link BDP:

BDP [Kbytes] = (link BW [Kbytes/sec] * Round-trip latency [Sec])

When multiple links 1..N are the links for which the WAE is optimizing traffic, the maximum BDP should be calculated as follows:

MaxBDP = Max (BDP(link 1),..,BDP(link N))

If the calculated MaxBDP is greater than the DefaultBDP for your WAE model, the Acceleration TCP settings should be modified to accommodate that calculated BDP.

Once you calculate the size of the Max BDP, enter a value that is equal to or greater than twice the Max BDP in the Send Buffer Size and Receive Buffer Size for the optimized connection on the Acceleration TCP Settings window.

Note These manually configured buffer sizes apply only if TCP adaptive buffering is disabled. TCP adaptive buffering is normally enabled, and allows the WAAS system to dynamically vary the buffer sizes. For more information on TCP adaptive buffering, see the "Modifying the TCP Adaptive Buffering Settings" section on page 13-63.

Modifying the TCP Adaptive Buffering Settings

In most cases, you do not need to modify the acceleration TCP adaptive buffering settings because your WAAS system automatically configures the TCP adaptive buffering settings based on the network bandwidth and delay experienced by each connection. Adaptive buffering allows the WAAS software to dynamically vary the size of the send and receive buffers to increase performance and more efficiently use the available network bandwidth.

To modify the acceleration TCP adaptive buffering settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Acceleration > TCP Adaptive Buffering Settings**. The TCP Adaptive Buffering Settings window appears.

Step 3 To enable TCP adaptive buffering, check the **Enable** check box. The default is enabled.

Step 4 In the Send Buffer Size and Receive Buffer Size fields, enter the maximum size in kilobytes of the send and receive buffers.

Step 5 Click **Submit**.

To configure the TCP adaptive buffer settings from the CLI, use the **tfo tcp adaptive-buffer-sizing** global configuration command:

```
WAE(config)# tfo tcp adaptive-buffer-sizing receive-buffer-max 8192
```

To disable TCP adaptive buffering from the CLI, use the **no tfo tcp adaptive-buffer-sizing enable** global configuration command.

To show the default and configured adaptive buffer sizes, use the **show tfo tcp** EXEC command.

CHAPTER 14

Configuring Virtual Blades

This chapter describes how to configure virtual blades, which are computer emulators that reside in a WAE or WAVE device. A virtual blade allows you to allocate WAE system resources for use by additional operating systems that you install on the WAE hardware. You can host third-party applications in the isolated environment provided by a virtual blade. For example, you could configure a virtual blade in a WAE device to run Windows printing and domain lookup services.

For detailed information on installing and configuring Windows on a virtual blade, see the *Cisco WAAS Installation and Configuration Guide for Windows on a Virtual Blade*.

Virtual blades are supported on all WAVE devices and on certain models of WAE devices. Virtual blades are not supported on WAVE devices operating as AppNav Controllers. On unsupported WAE devices, the virtual blade configuration screens are nonfunctional.

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers, WAEs, and WAVEs in your network. The term WAE refers to WAE and WAVE appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

This chapter contains the following sections:

About Virtual Blades

Virtual blades act as computer emulators within your WAAS device. You can install a guest operating system and applications on the virtual blade to work with your WAAS system and provide additional services for the users on your network.

Note WAAS virtual blades can host Windows Server 2003, 2008, and 2008 R2 operating systems with a broad range of services including, but not limited to, Windows services like Active Directory, Print Services, DHCP, SCCM, and DNS services, plus third-party and custom developed applications.

Each virtual blade has its own virtualized CPUs, memory, firmware, disk drives, CD drives, and network interface cards. A virtual host bridge controls communications between the virtual blade, your WAAS device, and the rest of your WAAS network.

Note When you configure a virtual blade on your WAAS device, system resources are reserved for the virtual blade. These resources are not available to your WAAS system, even if the virtual blade is not active. This can affect the performance of your WAAS system.

Each virtual blade includes a Virtual Network Computing (VNC) server that allows you to use a VNC client to connect to the virtual blade console so that you can observe and manage the guest operating system. The VNC client will need the IP address of the virtual blade console, which is the IP address of the WAAS device with the virtual blade number specified after a colon (for example: 10.10.10.40:1).

Note The VNC client adds 5900 to the virtual blade number to determine the port to connect to on the virtual blade. For virtual blade 1 this would be port 5901. An alternate way of specifying the port number is to specify the IP address followed by a space and then the actual port number, for example: 10.10.10.40 5901.

With a virtual blade you can perform the following activities:

- Configure system characteristics of the virtual blade environment
- Install an operating system and applications
- Configure the network flow to and from the virtual blade
- Start and stop the virtual blade

Table 14-1 lists an overview of the steps required to set up and enable one or more virtual blades on your WAAS device.

Table 14-1 Virtual Blade Configuration Overview

Step	Description
1. Prepare your WAE for using virtual blades.	Enables the virtual blade feature on your WAE-674 (not necessary on other platforms). See the "Preparing to Use Virtual Blades" section on page 14-3.
2. Configure the virtual blade system parameters.	Sets up system resources and interfaces for your virtual blade. See the "Configuring Virtual Blades" section on page 14-4.
3. Start the virtual blade on your WAE.	Starts the virtual blade running. See the "Enabling and Disabling Virtual Blades" section on page 14-8.
4. Transfer files to the virtual blade.	Copies files to the WAAS device hard drive for use by your virtual blade. See the "Copying a Disk Image to a Virtual Blade" section on page 14-10, and see the "Backing Up and Restoring a Virtual Blade" section on page 14-11.

Preparing to Use Virtual Blades

Note This procedure applies only to WAE-674 devices. Virtual blades are always enabled on WAVE platforms. You cannot disable a virtual blade on a WAVE device.

Before you configure and enable a virtual blade on your WAE-674 device, follow these steps:

Step 1 Ensure that the Virtual-Blade license is installed on the device. For more information, see the "Managing Software Licenses" section on page 10-3.

Step 2 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Note You can enable and configure virtual blades only on application accelerators, not on Central Manager devices or ANCs, and only on individual WAAS devices. You cannot configure virtual blades on device groups.

Step 3 Choose **Admin > Virtualization > General Settings**. The General Settings window appears.

Step 4 Check **Enable Virtualization** to enable virtualization.

Note If virtualization is enabled, it can be disabled only by reinstalling WAAS using the recovery CD.

Step 5 Click **Submit**.

You are prompted to confirm that you want to modify general settings. Doing so will reboot the WAE. After the reboot, the WAE will have a disk partition and other resources reserved for virtual blade use.

Note You will not be able to undo this change unless you restore the WAE from the Recovery CD.

Note When you configure a virtual blade on your WAE device, system resources are reserved for the virtual blade. These resources are not available to your WAAS system, even if the virtual blade is not active. This can affect the performance of your WAAS system.

Step 6 Click **OK**. The WAE restarts.

Step 7 Locate the disk or image of the operating system that you want to run on the virtual blade. Make sure that you either have the CD-ROM available, or that you have copied the disk image to the WAE hard drive. See the "Copying a Disk Image to a Virtual Blade" section on page 14-10.

To enable virtualization with the WAAS CLI, use the **virtual-blade** global configuration command.

Configuring Virtual Blades

This section describes how to configure a new virtual blade or edit an existing blade. You can configure resources such as the virtual blade number, description, boot method, disk allocation, CPU list, and other parameters. Note that after a virtual blade is initially configured, the only resource parameters that can be changed are memory and the bridged interface. To change these parameters on a virtual blade, stop the virtual blade first, then start the virtual blade after making changes.

To configure a virtual blade on your WAAS device, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Admin > Virtualization > Virtual Blades**. The Virtual Blade Entries window appears (see Figure 14-1).

Figure 14-1 Virtual Blade Entries List Window

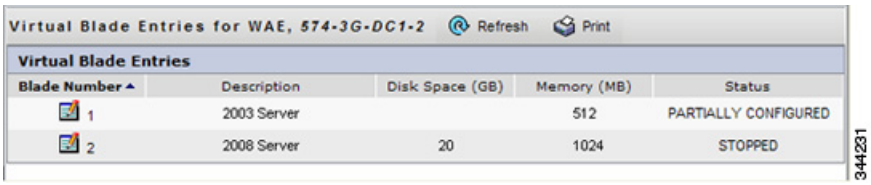

Any existing virtual blades are displayed in the Virtual Blade Entries list.

Note Status is not available for virtual blades running on WAAS version 4.1.1, so the status column shows NOT AVAILABLE. To get status, see the virtual blade Actions window (see Figure 14-3 on page 14-9).

Step 3 Click the **Edit** icon next to the virtual blade that you want to configure, or click the **Create** button to create a new virtual blade. The Virtual Blade configuration window appears (see Figure 14-2).

Figure 14-2 Virtual Blade Configuration Window

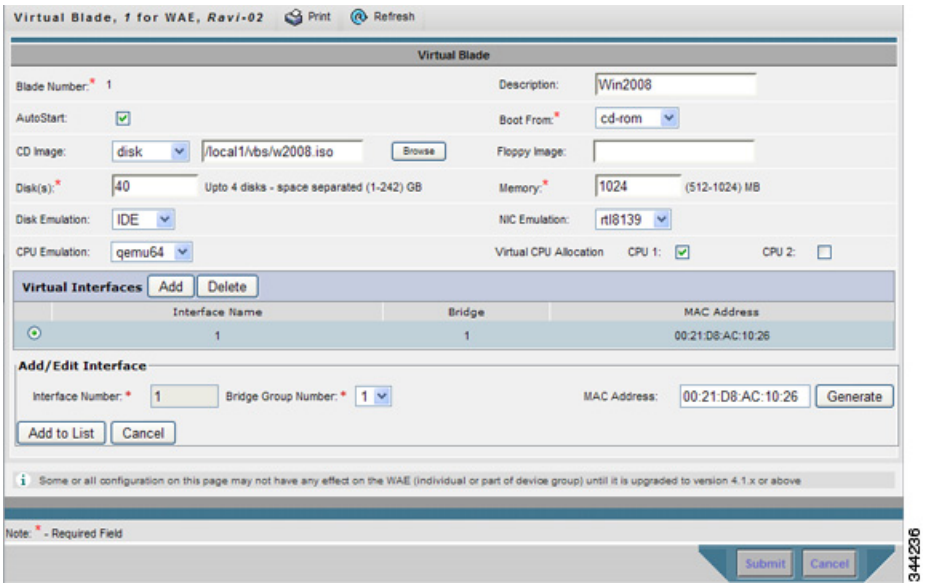

Step 4 Configure the virtual blade system parameters as needed to run your operating system and applications:

a. If you are creating a new virtual blade, type the number of the virtual blade that you want to create in the Blade Number field.

The number of virtual blades you can configure depends on which model of WAAS appliance you are using, and the amount of memory installed in that appliance.

b. (Optional) In the Description field, type a brief description of the virtual blade.

c. (Optional) Check **Autostart** to set the virtual blade to start automatically whenever the WAE starts.

d. Choose a source for the virtual blade to boot from by using the Boot From list, as follows:

- Choose **cd-rom** (the default) to boot the virtual blade from a physical CD or a CD image (.iso image file stored in the /local1/vbs directory). Use this choice before you have installed a guest OS, to boot from a guest OS installer CD.
- Choose **disk** to boot the virtual blade from a guest OS installed on the WAAS device hard drive. Use this choice after you have installed a guest OS, to boot from the installed guest OS.
- Choose **network** to boot the virtual blade from a network location (requires PXE to be enabled on your network). Use this choice to install the same version of software to many virtual blades, or to boot each virtual blade with a complete OS stored and managed in a centralized network location.

If you are network booting to install the guest OS, you may want to configure the virtual blade to boot from the disk on subsequent boots. You can do this by modifying the **boot from** parameter while the virtual blade is running.

e. Specify the location of a CD image by using the CD Image list. If you specified **cd-rom** for the Boot From list, the CD Image setting is required and configures the location of the boot image. If you specified **disk** or **network** for the Boot From list, the CD Image setting is optional and configures the location of a CD-ROM image that is made available to the guest OS (but is not used for booting). The choices for the CD Image are as follows:

- Choose **cd-rom** to read the CD image from a physical CD in the WAAS device CD-ROM drive.

– Choose **disk** to read the CD image from an ISO file on the WAAS device hard drive. If you choose **disk**, click the **Browse** button and select the ISO file from the /local1/vbs directory. The **Browse** button is shown only if there are files in the /local1/vbs directory. If you need to copy an ISO file to the /local1/vbs directory, see the "Copying a Disk Image to a Virtual Blade" section on page 14-10.

The CD image can be changed during operation, from the Virtual Blade Actions page, by clicking **Eject CD-ROM**, followed by clicking **Use CD-ROM** or specifying an ISO disk image and clicking **Set Image**.

f. If you want to reserve resources for a virtual floppy disk on your virtual blade, enter the pathname of the floppy disk image in the Floppy Image field. The path must be /local1/vbs/*filename*.

g. In the Disk Space field, enter the size of the virtual hard disk, in Gigabytes, that you want to allocate for the virtual blade.

You can configure up to four virtual hard disks on the virtual blade by specifying four hard disk sizes, separated by spaces, as shown in Figure 14-2. If you are using IDE disk emulation, you must specify 0 for the size of the third disk, because this IDE bus position is used for a CD-ROM.

Caution Do not use this Central Manager window to manage a virtual blade on a WAAS device that is running WAAS version 4.1.1 if multiple virtual hard disks are configured on the virtual blade. If you use this Central Manager window to change any part of the virtual blade configuration of a WAAS device running WAAS version 4.1.1 that already has multiple virtual hard disks configured, the Central Manager removes the disk configuration for all disks after the first disk and erases the other virtual disks.

You cannot configure multiple virtual hard disks from the Central Manager for WAAS devices that are running WAAS version 4.1.1. Instead, use the **disk** virtual blade configuration command from the CLI on the WAE.

h. In the Memory field, allocate the amount of WAE memory, in Megabytes, that you want to make available for the virtual blade.

The amount of memory that can be allocated for a virtual blade depends on the amount of memory in your WAE or WAVE appliance, and on the amount of memory that is assigned to other virtual blades. The minimum amount of memory that you can allocate for a single virtual blade is 512 MB.

i. In the Disk Emulation list choose the type of disk emulation that the virtual blade uses. Choose **IDE**.

IDE specifies an IDE (ATA) type disk emulator. **Virtio** specifies a generic disk controller emulator optimized for virtual machines.

Note If you select the virtio emulator, you must have the paravirtualization (PV) drivers installed on your system and you must validate this configuration. Virtio for disk emulation is provided on an experimental basis only and is not a supported feature.

j. On the NIC Emulation list, choose the type of network interface card emulation that the virtual blade uses. Choose **rtl8139**, **E1000**, or **virtio**.

Rtl8139 specifies a Realtek network card emulator, E1000 specifies an Intel PRO/1000 network card emulator, and virtio specifies a generic NIC emulator optimized for virtual machines. If you choose the virtio emulator, you must have the paravirtualization (PV) drivers installed on your system. (See the "Installing Paravirtualization Drivers" section on page 14-8.)

k. On the CPU Emulation list, choose the type of CPU emulation that the virtual blade uses. Choose **qemu64** (for a 64-bit processor emulator) or **qemu32** (for a 32-bit processor emulator).

l. In the Virtual CPU Allocation field, choose each CPU to assign to the virtual blade.

If you choose multiple CPUs, the CPUs are used in SMP mode. If two CPUs are available, by default odd numbered virtual blades use CPU 1, and even numbered virtual blades use CPU 2. If four CPUs are available, by default virtual blades are distributed among the CPUs; virtual blades 1 through 4 use CPUs 1 through 4, respectively, and virtual blades 5 and 6 again use CPUs 1 and 2, respectively.

You may configure any combination of CPUs; however enabling a virtual blade to use more than one core in SMP mode may interfere with another virtual blade using the same core. In this case, a warning appears.

Note A running virtual blade can be moved between CPUs, but the virtual blade must be stopped to add or remove CPUs.

The number of CPUs available for virtual blades depends on the device. Half the CPUs on a device are reserved for the WAAS software. If no virtual blades are started, all CPUs are used for the WAAS software.

Step 5 Configure the bridge that you want to use between the virtual blade and the physical interfaces on your WAE by doing the following:

a. Configure a bridge group and a bridge virtual interface, and add a physical interface to the bridge group as described in the "Bridging to a Virtual Blade Interface" section on page 6-17.

b. In the Virtual Interfaces pane, click the **Add** button.

c. In the lower part of the Virtual Interfaces pane, under Add/Edit Interface, in the **Interface Number** field enter the virtual blade interface to be bridged. Valid values are 1 or 2.

d. In the Bridge Group Number drop-down list, choose the bridge group in which to place the virtual blade interface.

e. In the MAC Address field, enter a MAC address for the bridged interface or click **Generate** to have WAAS generate the MAC address for you.

f. Click **Add to List** to add the virtual interface to the Virtual Interfaces list.

Step 6 In the Virtual Interfaces pane, click the radio button of the virtual interface that you want to use.

Step 7 Click **Submit**.

Note To access the virtual blade console use the IP address of the bridge virtual interface, with the virtual blade number specified as the port number (separated by a colon). For example if you bridged to the BVI interface with an IP address of 10.10.10.20, use **10.10.10.20:1** to get to the virtual blade 1 console.

To configure virtual blades with the WAAS CLI, use the following commands:

- **virtual-blade** (to enter virtual blade configuration mode)
- **(config-vb) autostart** to enable autostart
- **(config-vb) boot** to set the boot device
- **(config-vb) cpu-list** to configure the CPU list
- **(config-vb) description** to enter a description for the virtual blade
- **(config-vb) device** to define the CPU, NIC, and disk emulators

- **(config-vb) disk** to allocate disk space for the virtual blade
- **(config-vb) interface** to bridge a virtual blade interface to a bridge group
- **(config-vb) memory** to allocate system memory for the virtual blade
- **(config-vb) vnc** to disable the VNC server on the virtual blade (it is enabled by default)

Installing Paravirtualization Drivers

To install paravirtualization drivers, perform the following steps:

Step 1 Download the paravirtualization drivers file (virtio-drivers.iso) from the Cisco.com Software Download website in the WAAS Tools area.

The VirtIO network drivers are availabe for the following Windows operating systems:

- Windows 2003 32-bit and 64-bit
- Windows Server 2008 32-bit and 64-bit
- Windows Server 2008 R2 64-bit

Step 2 Copy the virtio-drivers.iso to the /vbs directory of your WAAS device using the following command:

```
wae# copy ftp disk ip_address source_dir virtio-drivers.iso vbs/virtio-drivers.iso
```

Where *ip_address* and *source_dir* are the IP address and source directory of the FTP server.

Step 3 Load the virtio-drivers.iso file on the virtual blade using the following command:

```
wae# virtual-blade 1 cd disk vbs/virtio-drivers.iso
```

Step 4 Choose one of the following driver methods for the installation:

- To install the driver for 2008 or 2008 R2 64-bit, execute the following command from a Windows command shell:

  ```
  c:\pnputil -i -a d:\inf\amd64\Win2008\netkvm.inf
  ```

- To install the driver for 2008 32-bit, execute the following command from a Windows command shell:

  ```
  c:\pnputil -i -a d:\inf\i386\Win2008\netkvm.inf
  ```

- To install the driver for Windows 2003, open Windows explorer and change the directory to one of the following:
  - d:\inf\i386\Win2003
  - d:\inf\amd64\Win2003

  Then right click netkvm.inf and select **Install** from the pop-up context menu.

Enabling and Disabling Virtual Blades

To enable or disable a virtual blade on your WAE, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Admin > Virtualization > Actions**. The Virtual Blade Actions window appears (see Figure 14-3).

Figure 14-3 Virtual Blade Actions Window

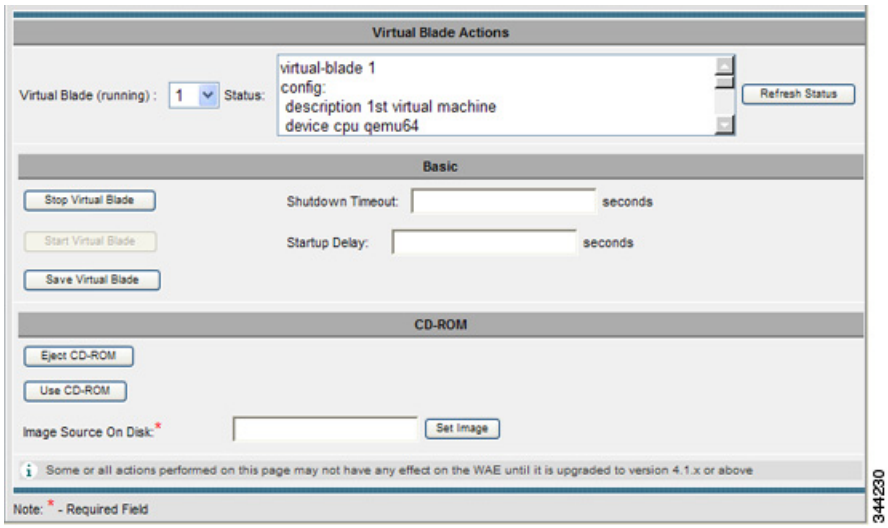

Step 3 In the Virtual Blade list, choose the virtual blade that you want to enable or disable. The status of the virtual blade is displayed in the Status field.

The default selection for the Virtual Blade list is All. When All is selected, the Status field displays the current status for all virtual blades.

Step 4 Click **Start Virtual Blade** to enable the selected virtual blade.

- (Optional) Enter a startup delay by typing a value, in seconds, in the **Startup Delay** field.

  The startup delay can be used to give you time to connect a VNC session to the console before the virtual blade boots, so you can observe the initial startup.

Step 5 Click **Stop Virtual Blade** to disable the selected virtual blade.

- (Optional) To give the virtual blade operating system time to shut down the virtual blade after you click the Stop Virtual Blade button, enter a value (in seconds) in the Shutdown Timeout field.

  The shutdown timeout provides a delay period during which the operating system can shutdown gracefully. If the operating system has not shut down the virtual blade by the end of this period, WAAS cancels the shutdown.

  If you set the Shutdown Timeout to 0, WAAS forces a shutdown immediately. A forced shutdown is comparable to pulling the power cord on a real computer.

  To avoid losing data in open programs running on the virtual blade, it is safer to have the operating system perform the shutdown.

Step 6 Click **Refresh Status** to refresh the status of the virtual blade after you make a change.

During virtual blade operation, the CD image can be changed by clicking **Eject CD-ROM**, followed by clicking **Use CD-ROM** (for a physical CD) or specifying an ISO disk image and clicking **Set Image**.

Note The operating system on your virtual blade does not shut down and restart when you reboot a WAAS device. When you reboot a WAE or WAVE device, the WAAS software saves the virtual blade in its current state and then restores that state when the reboot is complete.

To enable a virtual blade with the WAAS CLI, use the **virtual-blade** *n* **start** EXEC command. To disable a virtual blade, use the **virtual-blade** *n* **stop** EXEC command.

To eject a CD or virtual CD image, use the **virtual-blade** *n* **cd eject** EXEC command.

To use a new CD that has been inserted in the CD-ROM drive, use the **virtual-blade** *n* **cd cd-rom** EXEC command.

To use a new CD ISO image from the WAE /local1/vbs directory, use the **virtual-blade** *n* **cd disk** *pathname* EXEC command.

Copying a Disk Image to a Virtual Blade

If you want to boot from a disk image stored on the WAAS device hard drive, you must copy that image file to the virtual blade staging area under the following directory: /local1/vbs.

To copy a disk image file to the /local1/vbs directory on your WAE, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Admin > Virtualization > File Transfers**. The Virtual Blade File Transfers window appears (see Figure 14-4).

Figure 14-4 Virtual Blade File Transfers Window

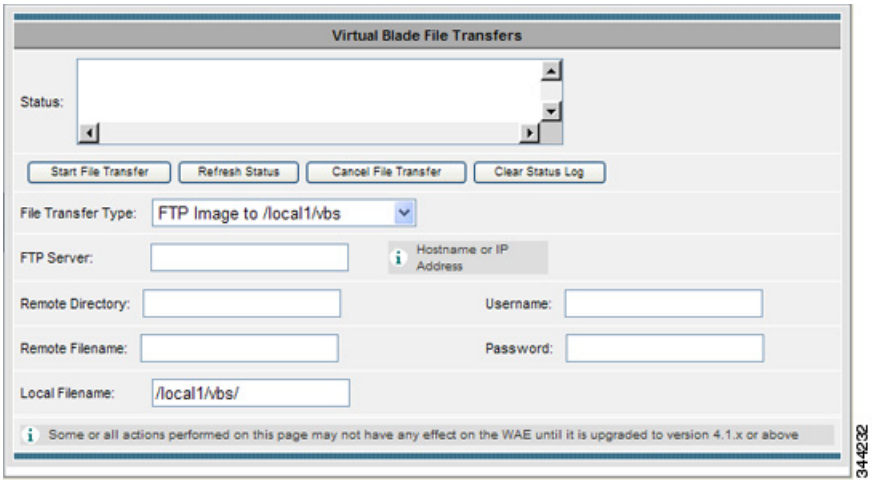

Step 3 In the File Transfer Type list, choose **FTP Image to /local1/vbs** (the default).

Step 4 In the FTP Server field, enter the IP address or hostname of the FTP server where the disk image resides.

Step 5 In the Remote Directory field, enter the path to the directory on the FTP server where the disk image resides.

Step 6 In the Remote Filename field, enter the filename of the disk image.

Step 7 In the Username and Password fields, enter a valid username and password for the FTP server.

Step 8 In the Local Filename field, enter the full path and filename where the disk image is to be stored on the WAE device. The directory path must be /local1/vbs/.

Step 9 Click **Start File Transfer** to start the file transfer.

File transfer status information is shown in the Status field. To refresh status information, click **Refresh Status**.

To cancel a file transfer, click **Cancel File Transfer**.

To clear the status information field, click **Clear Status Log**.

You can also back up and restore virtual blade disk images from this window. For details, see the "Backing Up and Restoring a Virtual Blade" section on page 14-11.

To copy the operating system ISO image to the virtual blade directory by using the CLI, use the **copy ftp disk** EXEC command. For example, the following command copies the boot image winserver.iso from the WAAS directory on FTP server 10.10.10.200 to the virtual blade directory on the WAE device (/local1/vbs/):

```
wae# copy ftp disk 10.10.10.200 WAAS winserver.iso /local1/vbs/winserver.iso
```

Backing Up and Restoring a Virtual Blade

You can back up and restore a virtual blade's disk image. The disk image is the bootable operating system and applications that run on the virtual blade. For example, your virtual blade might have a disk image of Windows Server 2003 running Print Services.

Note WAAS virtual blades support only the Windows Server 2003 or Window Server 2008 operating systems, and the Active Directory, Print Services, DHCP, and DNS services. Although other operating systems and applications may operate on a virtual blade, the WAAS virtual blade does not support these other operating systems and applications.

To back up a virtual blade disk image to an FTP server, follow these steps:

Step 1 Stop the virtual blade that you want to back up. To stop a virtual blade from the WAAS Central Manager, use the procedure described in the "Enabling and Disabling Virtual Blades" section on page 14-8.

Step 2 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 3 Choose **Admin > Virtualization > File Transfers**. The Virtual Blade File Transfers window appears (see Figure 14-4 on page 14-10).

Step 4 In the File Transfer Type list, choose **Backup Virtual Blade to FTP**.

Step 5 In the FTP Server field, enter the IP address or hostname of the FTP server to which you want to back up the virtual blade disk image.

Step 6 In the Remote Directory field, enter the path to the directory on the FTP server where you want to copy the disk image.

Step 7 In the Remote Filename field, enter the filename for the disk image.

Step 8 In the Username and Password fields, enter a valid username and password for the FTP server.

Step 9 In the Virtual Blade No. field, enter the number of the virtual blade that you want to back up.

Step 10 In the Disk No. field, enter the number of the virtual blade disk that you want to back up. For backing up a virtual blade running Microsoft Windows Server, always enter **1**.

Step 11 Click **Start File Transfer** to start the file transfer.

To restore a virtual blade disk image that you previously backed up, follow the procedure above but choose **Restore Virtual Blade from FTP** from the File Transfer Type list.

If a virtual blade is already configured on the WAE, you must remove it before you can restore a virtual blade disk image.

Note If the virtual blade configuration was removed before restoring the virtual blade disk image, you must reconfigure the virtual blade after the restore operation. Configure all virtual blade system parameters except for the disk sizes, which are configured by the restore operation.

File transfer status information is shown in the Status field. To refresh status information, click **Refresh Status**.

To cancel a file transfer, click **Cancel File Transfer**.

To clear the status information field, click **Clear Status Log**.

At a given time, only a single backup or restore operation can be in progress.

To back up the disk image of a virtual blade on your WAE to an FTP server by using the CLI, use the **copy virtual-blade** EXEC command. For example, the following command transfers the file file.img from disk 1 on virtual blade 1 to FTP server 10.75.16.234:

```
wae# copy virtual-blade 1 disk 1 ftp 10.75.16.234 / file.img
```

To restore a disk image to a virtual blade on your WAE, use the **copy ftp virtual-blade** EXEC command. For example, the following command transfers the file file.img from the FTP server 10.75.16.234 to disk 1 on virtual blade 1:

```
wae# copy ftp virtual-blade 1 disk 1 10.75.16.234 / file.img
```

CHAPTER 15

Configuring the Network Analysis Module

This chapter provides information about the integration of the Cisco Network Analysis Module (NAM) in the WAAS Central Manager and describes how to configure and use the NAM.

This chapter includes the following sections:

Information About NAM Integration

WAAS is enhanced with application performance monitoring capabilities when you integrate the WAAS Central Manager with the Cisco Network Analysis Module (NAM) Traffic Analyzer software.

The NAM Traffic Analyzer software enables network managers to understand, manage, and improve how applications and services are delivered to end users by combining flow-based and packet-based analysis into one solution. With the NAM, you can perform traffic analysis of applications, hosts, and conversations, make performance-based measurements on application, server, and network latency, and use Quality of Service (QOS) metrics for network-based services and problem analysis using packet captures. The Cisco NAM includes an embedded, web-based Traffic Analyzer GUI that provides quick access to the configuration menus and presents easy-to-read performance monitoring and analysis on network traffic.

The architecture for WAAS Central Manager and NAM integration has the NAM deployed as a virtual blade in one of the data center WAEs. While this deployment eliminates the need for separate hardware in Proof of Concept (POCs) deployments, users are not limited to this deployment and can deploy NAM 5.1 in any form factor such as virtual blade, physical blade, or appliance.

Prerequisites

The NAM integration has the following prerequisites:

- The WAAS 4.4.1 or later Central Manager is installed and configured.
- The NAM 5.1 hardware and software are installed.

- The following configurations are performed:
  - Enable HTTP or HTTPS
  - Create an admin web user account.
  - Create a MonitorView web user account.
- Both the WAAS Central Manager and the client computer from which you connect to the Central Manager must be able to access the configured NAM server on the network.

For more information, see the *Cisco Network Analysis Module Installation and Configuration Guide*.

To install the NAM Virtual Service Blade (VSB) on a WAAS appliance, configure the NAM VSB, and generate useful reports to demonstrate the impact of WAAS in a Proof of Concept (POC), see the *Cisco Prime Network Analysis Module for WAAS Virtual Blade Installation and Configuration Guide, 5.1*.

Guidelines and Limitations

The NAM integration feature has the following configuration guidelines and limitations:

Supported Deployments

In this release, the following types of deployments are supported:

- POC deployments
- Small and medium production networks that can be monitored by one instance of NAM (virtual blade, physical blade, or appliance). In this release, only one NAM instance is supported, which means that large enterprises that require more than one NAM instance to handle their network capacity must be managed separately without the WAAS-CM integration.

Limitations

- Browser limitations—Certain browser settings can limit the functionality of the NAM integration. For example, if Internet Explorer privacy settings are set to the default, Medium, the integration does not work because of cookie restrictions. Specify the privacy settings as Low.
- Printing limitations - When you print the NAM windows in a PDF format, they do not produce the desired output.
- When duplicate data is reported by multiple WAE data sources, the NAM does not automatically remove duplication of this data. Use the Data Source selector in the dashboards and charts to address this limitation.

Configuring the NAM

This section includes the following topics:

Task Flow for Configuring the NAM

This section includes the following topics:

Basic Configuration

The basic NAM configuration includes the following tasks:

- Configuring the setup (see Configuring the Basic Setup, page 15-3).
  - Connect to a NAM server by providing the server's IP address, protocol, and port.
  - Establish account credentials.
  - Associate a WAAS device group or WAAS Express device group with configured policies.
  - Enable Flow Agent.
- Configuring Sites—To display accurate data on charts and dashboards, every site on which WAAS is planned to be deployed must be configured on the NAM (see Configuring a Site, page 15-5).

  Configuring sites involves the following tasks:
  - Define sites
  - Use definition rules
  - Specify sites using subnets
- Configuring monitored servers (see Configuring a WAAS Monitored Server, page 15-10).
  - Specify the servers to be monitored by the NAM using the WAAS device's flow monitoring.
  - Enabling NetFlow and flow agent data sources on the actual devices, with the NAM as the collector, to automatically create the device entries in the NAM.

Advanced Configuration

Advanced NAM configuration includes the following tasks:

- Configuring and synchronizing user-defined Classifiers and Applications with the NAM (see Synchronizing Classifiers and Applications, page 15-10).
- Creating and editing an auto-created WAAS data source to monitor WAAS traffic statistics (see Configuring a Data Source, page 15-11).
- Changing system preferences (see Setting Preferences, page 15-14).
- Launching the NAM user interface (see Launching the NAM User Interface, page 15-14).

Configuring the Basic Setup

Only device group level policy configurations are applicable for NAM.

Step 1 From the WAAS Central manager menu, choose **Configure > Network Analysis Module > Basics > Setup**.

The Setup window appears. This window allows you to configure the NAM IP address and accounts.

Figure 15-1 Setup Window

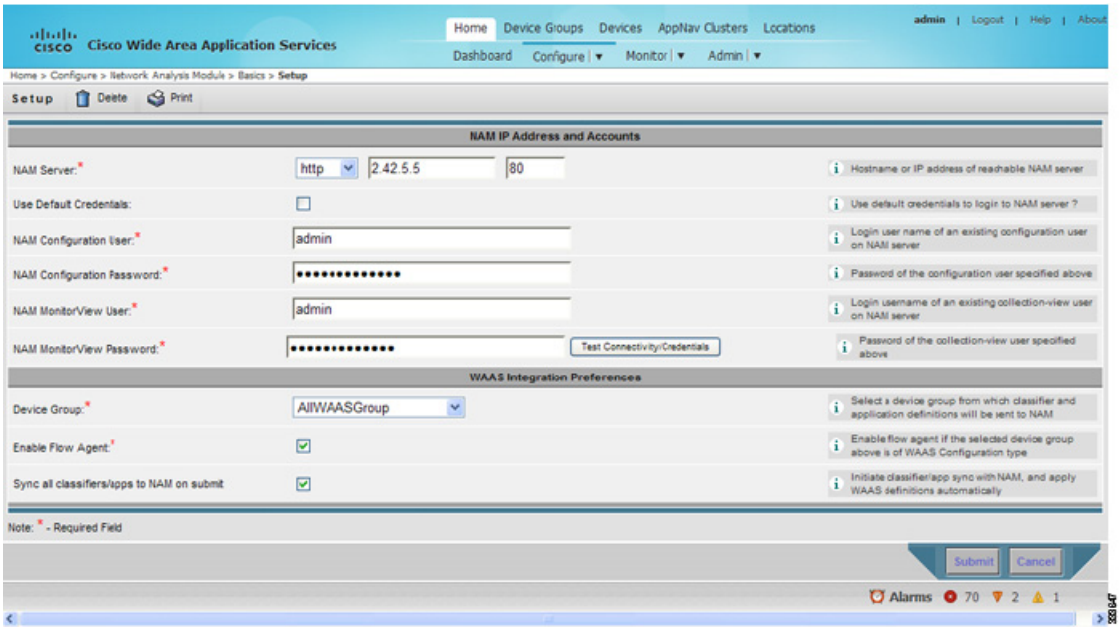

Step 2 In the NAM Server area, provide the following information:

- Choose either HTTP or HTTPS depending on the access that was configured during the installation of NAM.
- Enter the hostname of the NAM server.
- Enter the IP address of the NAM server.

To set up a sites or sites on the NAM module, perform the following steps:

Step 3 To use the pre-configured login credentials to access the NAM server, select the **Use Default credentials** option. Proceed to Step 8.

The following preconfigured login credentials are specified in the Central Manager:

- Configuration user:
  - Username—admin
  - Password—admin
- MonitorView user:
  - Username—waasro
  - Password—waasrao

These credentials work only if you configured the NAM with them explicitly after installation or you installed the NAM virtual blade with these credentials for POC deployments.

Step 4 In the NAM Configuration User field, enter the username of an existing configuration user on the NAM server.

Step 5 In the NAM Configuration Password field, enter the password of the configuration user that was specified in Step 3.

Step 6 In the NAM MonitorView User field, enter the username of an existing collection-view user configured on the NAM server.

Step 7 In the NAM MonitorView Password field, enter the password of an existing collection-view user that you specified in Step 3.

Step 8 Click the **Test Connectivity/Credentials** button, to test if the NAM server is accessible and to check if the user credentials that you specified are valid.

Step 9 The WAAS Integration Preferences area allows you to configure a WAAS device group to work with the NAM server.

a. From the Device Group drop-down list, choose a device group from which WAAS applications and classifier definitions are pushed to the NAM when performing a synchronization operation.

The AllWAASDevices or AllWAASExpressDevices device group is the default selection for POC deployments. For production deployments, choose a suitable device group with a subset of devices for which you require the NAM integration and APM functionality.

b. Choose the **Enable Flow Agent** option to enable sending of flow agent reports from the WAAS devices on the selected device group to NAM.

This option is disabled for WAAS Express device group because WAAS Express does not support the flow agent/flow monitor. In this scenario, you must use a NAM Performance Agent (PA) from Cisco IOS routers to view the response time metrics. The NAM charts that display response times in the WAAS Central Manager also automatically handle the PA from routers.

c. Choose the **Sync all classifiers/apps to NAM on submit** option to initiate a classifier and application synchronization with NAM and to apply WAAS definitions automatically.

Step 10 Click the **Submit** button.

Configuring a Site

A site is a collection of hosts, or network endpoints, partitioned into views that helps you to monitor traffic and to troubleshoot problems. These views allow you to see measurements of application performance on networks where WAAS devices are deployed and dashboards that show the traffic levels between sites and alarms levels per site. You can use other NAM features without defining any sites (the default configuration).

If you have set up sites, you can choose a particular site to view in the Interactive Report and view data relevant to that site only. In some cases, you can select both a client site and a server site to view data that pertains to the interaction between hosts at different sites.

Note If you configure multiple data sources for the same site, the same traffic might be accounted for more than once, which results in inflated traffic statistics. For example, if you configure the NAM to receive SPAN traffic for a particular site, and it is also receiving NetFlow records for that same site, both SPAN traffic and NetFlow records are combined into the traffic statistics. In this case, if you want to see only the statistics for a particular data source, you need to use the Interactive Report window on the left side of the window to specify both the site and data source.

Note Classification of received data from data sources to sites is done only after the sites are configured. Any old data from these data sources before the sites were configured are counted under the default 'Unassigned' site.

The site definition is very flexible and can accommodate various scenarios. The site definition is used not only for viewing of data but for data export and data retention as well. Typically, a site is defined by its subnet(s), but a site can also be defined using the following rules:

- Subnet (IP address prefix)
- Subnet from a data source
- Subnet from a given VLAN of a SPAN data source
- WAE device serving the site

We recommend that you define sites using subnets whenever possible.

Note The same rule cannot be defined in multiple sites.

Note If you are configuring a WAAS device, you must add WAAS servers to the NAM. See Auto Creating a New WAAS Device.

To display accurate data on charts and dashboards, you must configure every site on which WAAS is planned to be deployed on the NAM. To get a breakdown of the traffic volume and response time for each branch and the data center, configure the IP subnets for all the sites that has WAAS deployed.

This section includes the following topics:

Definition Rules

Typically, subnets alone are sufficient to define a site. For example:

```
Site Data-Center = subnet 172.20.0.0/16
```

In certain scenarios when there are overlapping IP address spaces in the networks (for example, in private networks where hosts from different sites have the same IP addresses), you can use data sources or VLANs to differentiate the subnets. For example:

```
Site NewYork = subnet 10.11.0.0/16 from "NDE-NewYork" data source.
Site LosAngeles = subnet 10.11.0.0/16 from "NDE-LosAngeles" data source.
Site Sale-Dept = subnet 10.11.0.0/16 from VLAN 10 of "DATA PORT 1" data source.
Site Finance-Dept = subnet 10.11.0.0/16 from VLAN 12 of "DATA PORT 1" data source.
```

This section includes the following topics:

- Specifying a Site Using Multiple Rules, page 15-7
- Resolving Ambiguity (Overlapping Site Definitions), page 15-7

Specifying a Site Using WAE devices (WAAS Data Sources)

For WAAS traffic, you can define a site associated with a WAE device without specifying the site's subnets. Simply select all of the WAAS data sources coming from the WAE device(s) serving that site.

```
Site SanJose = WAE-SJ-Client, WAE-SJ-CltWAN, and WAE-SJ-Passthrough data sources.
```

Note Cisco recommends that you use subnets to specify WAAS-optimized sites. Use this method only if the site's subnets cannot be determined.

Specifying a Site Using Multiple Rules

You can define a site using a combination of multiple rules described above. For example, if a site has both optimized and nonoptimized traffic, it can be defined using a combination of WAAS data sources and a subnet from a NetFlow Data Export (NDE) data source.

When you define a site using multiple data sources, be careful to make sure that those data sources do not have duplicated traffic to avoid counting the site traffic statistics twice.

Resolving Ambiguity (Overlapping Site Definitions)

Conflicting rules are not allowed in site definitions. Of the following two scenarios, the second one is not allowed:

```
1.2.3.0/24 from SPAN1 = SiteA
1.2.3.0/24 from SPAN1 = SiteB
```

Using a prefix is the preferred method. The data source and VLAN are secondary. In the following two scenarios, the first receives the higher priority:

```
1.2.3.0/24 = Site D
WAE1-Client datasrc = Site E
```

The longest prefix has higher priority. It has the same data source/VLAN. In the following two scenarios, the first receives the higher priority:

```
1.2.3.0/24 from SPAN1 = Site A
1.2.0.0/16 from SPAN1 = Site C
```

The more refined (specific) rule has higher priority. In the following two scenarios, the first would receive the higher priority.

```
1.2.3.0/24 from SPAN1 = Site A
1.2.3.0/24 (any datasrc) = Site D
```

Viewing Defined Sites

You can view a defined site.

Step 1 From the WAAS Central manager menu, choose **Configure > Network Analysis Module > Basics > Sites**.

The Sites window appears. Defined sites are listed in the table.

The following details for the sites display:

- Name—Lists the name of the site.
- Description—Describes what the site includes.
- Rule—Lists the first rule that is assigned to the selected site. If you see periods next to the site rule (...), that means that multiple rules were created for that site. To see the list of all rules, click the quick view icon (after highlighting the site, click the small arrow on the right).
- Status—Shows if the site is enabled or disabled.

Defining Sites

You can define a site.

Step 1 From the WAAS Central Manager menu, choose **Configure > Network Analysis Module > Basics > Sites**.

The Sites window appears. This window lists the sites that are set up on the NAM module.

Step 2 Click the **Create** button. The Sites Configuration window displays.

Step 3 In the Name field, enter a name for the site.

Step 4 In the Description field, enter a description for the site.

Step 5 Check the **Disable Sites** check box if you want the NAM to skip this site when classifying traffic. This feature is useful if the site is no longer active, but you would still like to access historical site data in the database. Otherwise, you should delete sites that are not needed.

Step 6 In the Subnet field, enter the IP address subnet (IPv4/IPv6 address and mask); for example, 10.1.1.0/24.

Step 7 Click the blue **i** to get information about the site rules.

Step 8 Click the **Detect** button to tell the NAM to look for subnets in the traffic. See Detecting a Subnet, page 15-8.

Step 9 In the Data Source field, specify the data source where the site traffic is coming from.

Leave this field blank if the site traffic can come from multiple data sources.

Step 10 In the VLAN field, specify the VLAN where the site traffic is coming from. This field is not valid for NDE and WAAS data sources.

Leave this field blank if the site traffic can come from multiple VLANs.

Step 11 Click the **Submit** button.

Note The “Unassigned” site (with a description of “Unclassified hosts”) includes any sites that do not match any of your site configurations. Sites are classified at the time the packets are processed.

Detecting a Subnet

You can detect a subnet.

Step 1 When you click the **Detect** button at **Configure > Network Analysis Module > Basics > Sites > Sites Configuration**, the NAM looks for subnets detected within in the past hour. The Subnet Configuration window displays. This window allows you specify the details of the sources in which you like NAM to detect subnets.

Step 2 In the Subnet Mask field, enter the subnet mask.

Note If the bit mask is less than 32, the NAM detects an IPv4 subnet. If the bit mask is between 32 and 64, the NAM detects an IPv6 subnet.

Step 3 From the Data Source drop-down list, choose the data source in which you would like to detect subnets.

Step 4 From the Interface drop-down list, choose the interface in which you would like to detect subnets.

Step 5 In the Filter Subnets within Network field, enter an IPv4 or IPv6 address.

Step 6 Check the **Unassigned site** check box to include sites that do not match any of your site configurations. Sites are classified at the time of packet processing.

Step 7 Click the **Detect** button. The NAM finds those subnets that meet the criteria that you entered.

Editing a Site

You can edit sites that have been created. Note that the “Unassigned” site cannot be edited or deleted.

Step 1 Choose **Configure > Network Analysis Module > Basics > Sites**. A list of configured sites appears.

Step 2 Choose the site that you want to edit.

Step 3 Click the **Edit** button. The Site Configuration window displays.

Step 4 Edit the required field (Name, Priority, Data Sources, or Prefix/Mask).

Step 5 Click the **Submit** button.

Deleting a Site

You can delete sites that have been created. Note that the “Unassigned” site cannot be deleted.

Step 1 Choose **Configure > Network Analysis Module > Basics > Sites**. A list of configured sites appears.

Step 2 Choose the site that you want to delete.

Step 3 Click the **Delete** button.

Configuring a WAAS Monitored Server

WAAS monitored servers specify the servers from which WAAS devices export traffic flow data to the NAM monitors. To enable WAAS monitoring, you must list the servers to be monitored by the NAM using the WAAS device's flow monitoring.

Note The NAM is unable to monitor WAAS traffic until you set up WAAS monitored servers. The NAM displays the status of WAAS devices as pending until you set up WAAS monitored servers.

This section includes the following topics:

- Adding a WAAS Monitored Server, page 15-10
- Deleting a WAAS Monitored Server, page 15-10

Adding a WAAS Monitored Server

You can add a WAAS monitored server.

Step 1 Choose **Configure > Network Analysis Module > Basics > Monitored Servers**. The WAAS Servers window appears.

Step 2 Choose **Select All** to add all the servers or select the required servers from the list.

Step 3 Click the **Add** button.

Deleting a WAAS Monitored Server

To delete a WAAS monitored server data source:

Step 1 Choose **Configure > Network Analysis Module > Basics > Monitored Servers**.

The WAAS Servers window appears.

Step 2 Choose the monitored WAAS server to delete, and click the **Delete** button.

A confirmation dialog box displays to ensure you want to delete the selected WAAS monitored server.

Step 3 Click the **OK** button to delete the WAAS monitored server.

Synchronizing Classifiers and Applications

You can synchronize the WAAS classifier and application definitions with the application and application groups in the NAM. A classifier and an application in WAAS are equivalent to an application and application group respectively in the NAM. WAAS applications and classifier definitions from the device group specified during the setup configuration are matched with those in the NAM server that WAAS is connected to. WAAS classifiers can contain source and destination IP addresses while the NAM recognizes an application on the basis of port numbers. Hence, only the WAAS classifiers that contain port numbers are synchronized.

You can view the results of the synchronization by following these steps:

Step 1 Choose **Configure > Network Analysis Module > Advanced > Classifier/App Sync**. The **Classifier/App Sync Preferences** window appears.

The results are displayed under the following categories:

- Conflicting classifiers/applications—You can choose one or all the WAAS classifiers/applications for synchronization with the NAM. By default, all the classifier/applications are selected.
- NAM-only applications/application groups—Applications/application groups in the NAM are displayed. If required, you can manually add these definitions in WAAS at the device-group or device levels.

Step 2 To view differences in classifier definitions in WAAS and the NAM, click on the arrow next to **Classifier Definition Differences**.

Step 3 Choose the WAAS classifiers that you want to synchronize with the NAM applications and provide the required information to define the filter criteria.

Step 4 Click the **Go** button. The differences in the definitions are displayed.

Step 5 To view applications/application groups in the NAM, click on the arrow next to **NAM-Only Applications**. Information on the applications/application groups is displayed. If required, you can manually add these definitions in WAAS at the device-group or device levels.

Step 6 To refresh the Classifier/App Sync page, click the **Refresh** button.

Step 7 Click the **Submit** button to start the synchronization process.

Configuring a Data Source

Data sources are the source of traffic for the NAM Traffic Analyzer. Some examples are physical data ports of the NAM where you get SPAN data, a specific router or switch that sends NetFlow to the NAM, or a WAAS device segment that sends data to the NAM or ERSPAN and that goes to the NAM's management port.

A new feature in NAM 5.0 is the auto discovery of data sources, in which you can click **Auto Create** so that the NAM can automatically discover the data sources. You can see details such as the IP addresses of devices that send packets to the NAM and the time that the last NDE packet was received (in NAM 4.x, this feature was called Listening Mode).

Note If you have configured sites, you can assign data sources to that particular site. If you do assign data sources to a site, and you also configure the data sources, the two could overlap because sites can also be a primary "view" into data sources. If there is a mismatch between the two, you do not see any data.

Note We recommend that you configure a site using subnets instead of selecting a data source.

The following areas contain specific information about the types of data sources:

- SPAN
- ERSPAN
- VACL
- NetFlow

- WAAS

The NAM Data Sources window lists the data sources that are configured for that NAM module The fields are as follows:

- Device—DATA PORT if it is a local physical port or the IP address of the learned device.
- Type—The source of traffic for the NAM.
  - DATA PORT if it is a local physical port.
  - WAAS, ERSPAN, or NETFLOW if a data stream is exported from the router, switch, or WAE device.
- Activity—Shows the most recent activity.
- Status—ACTIVE or INACTIVE.
- Data Source—Name given to the data source.
- Data Source Details—Physical Port, or information about the data source being enabled or disabled.

This section includes the following topics:

- Adding a Data Source for a New WAAS Device, page 15-12
- Auto Creating a New WAAS Device, page 15-13
- Editing a WAAS Data Source, page 15-13
- Deleting a WAAS Data Source, page 15-13

Adding a Data Source for a New WAAS Device

The NAM uses WAAS data sources to monitor traffic that is collected from different WAAS segments: Client, Client WAN, Server WAN, Server, and Passthrough. Each WAAS segment is represented by a data source. You can set up the NAM to monitor and report other traffic statistics of the WAAS data sources such as application, host, and conversation information in addition to the monitored Response Time metrics.

Adding a WAAS device is not usually necessary because export-enabled WAAS devices are detected and added automatically.

To manually add a WAAS device to the list of devices monitored by the NAM:

Step 1 Choose **Configure > Network Analysis Module > Advanced > Data Sources**.

Step 2 Click the **Create** button.

The NAM Data Source Configuration dialog box appears.

Step 3 Choose **WAAS** from the list of Types.

Step 4 In the IP field, enter the device IP address.

Step 5 Check the check boxes for the appropriate WAAS segments.

You can configure the WAAS data sources to monitor the following WAAS segments:

- Client—Configures the WAE device to export the original (LAN side) TCP flows that originated from its clients to the NAM for monitoring.
- Client WAN—Configures the WAE device to export the optimized (WAN side) TCP flows originated from its clients to the NAM for monitoring.

- Server WAN—Configures the WAE device to export the optimized (WAN side) TCP flows from its servers to the NAM for monitoring.
- Server—Configures the WAE device to export the original (LAN side) TCP flows from its servers to the NAM for monitoring.
- Passthrough—This setting configures the WAE device to export the TCP flows that are passed through unoptimized.

Step 6 Click the **Submit** button to add the new WAAS custom data source.

Auto Creating a New WAAS Device

If you have numerous WAE devices, you can set up the NAM to configure newly discovered WAE devices using a predefined configuration template using the NAM Auto Config option.

Note If most of your WAE devices are edge WAE devices, you might want to set the auto config option as an edge device, and manually configure the data center WAE. For example, choose the Client segment for monitoring.

Step 1 Choose **Configure > Network Analysis Module > Advanced > Data Sources**. The data sources appear.

Step 2 Click the **Auto Create** button.

The NAM Data Source Configuration dialog box appears.

Step 3 Check the **WAAS** check box.

Step 4 Check the check boxes for the required segments. See Adding a Data Source for a New WAAS Device, page 15-12, for more information.

Step 5 Click the **Submit** button to add the new WAAS custom data source.

Editing a WAAS Data Source

You can edit a WAAS device's custom data source.

Step 1 Choose **Configure > Network Analysis Module > Advanced > Data Sources**. The data sources appear.

Step 2 Click the WAAS device that you want to modify, and click the **Edit** button. The NAM Data Source Configuration dialog box appears.

Step 3 Modify the segments as required.

Step 4 Click the **Edit** button to edit the WAAS custom data source.

Deleting a WAAS Data Source

You delete a WAAS custom data source.

Step 1 Choose **Configure > Network Analysis Module > Advanced > Data Sources**. The data sources appear.

Step 2 Choose the WAAS custom data source that you want to delete, and click the **Delete** button.

A confirmation dialog box appears to ensure that you want to delete the selected WAAS monitored server.

Step 3 Click **OK** to delete the WAAS custom data source.

Setting Preferences

You can configure characteristics such as NAM display, audit trail, and file format preferences for the NAM module.

Step 1 Choose **Configure > Network Analysis Module > Advanced > Preferences**. The Preferences window appears.

Step 2 Specify the following preferences:

- Refresh Interval (60-3600 sec)—Amount of time between the refresh of information on dashboards.
- Top N Entries (1-10—Number of colored bars on the Top N charts.
- Perform IP Host Name Resolution—Wherever an IP address appears, it gets translated to a hostname via a DNS lookup.
- Data Displayed In—Data displayed in Bytes or Bits.
- International Notation—Choose the way you would like the numbers to appear.
- Audit Trail—The Audit Trail option displays a listing of recent critical activities that have been recorded in an internal syslog log file. Syslog messages can also be sent to an external log.

Step 3 Click **Submit** to save your configurations

Launching the NAM User Interface

You can launch the NAM user interface to perform advanced configuration and monitoring tasks.

To launch the NAM user interface, do the following:

Choose **Configure > Network Analysis Module > Advanced > Launch NAM GUI**. A new window or a tab (depending on your browser settings) opens displaying a NAM session that uses the existing login credentials.

Monitoring and Analyzing Traffic

The monitoring and analyzing traffic feature provides intuitive workflows and interactive reporting capabilities.

The monitoring and analyzing dashboards allow you to view network traffic, application performance, site performance, and alarms at a glance.

This section provides information about monitoring your network traffic and analyzing the information presented.

This section contains the following topics:

Navigation

This section includes the following topics:

Interactive Report

On most monitoring dashboards, you can use the Interactive Report on the left to redefine the parameters of the information displayed in the dashboards. Click the **Filter** button to change the parameters of the information that appears in the charts.

You can choose from various parameters, such as the time interval for the data being displayed. An asterisk represents required fields.

The reporting time interval selection changes depending upon the dashboard that you are viewing, and the NAM platform that you are using:

- The NAM appliance supports the following short term intervals: Last 5 minutes, last 15 minutes, last 1 hour, last 4 hours, and last 8 hours.
- The Branch Routers (NME-NAM) support the following short term intervals: Last 5 minutes, last 15 minutes, and last 1 hour.
- The other platforms support the following short term intervals: Last 5 minutes, last 15 minutes, last 1 hour, and last 4 hours.
- The Long Term interval selections (Last 1 day, 1 week, and 1 month) are disabled from the following dashboards: RTP Streams, Voice Call Statistics, Calls Tables, RTP Conversations, Host Conversations, Conversations, and Response Time Details Views.
- A maximum interval for up to 1 hour is supported for the following dashboards: RTP Streams, Voice Call Statistics, Calls Tables, RTP Conversations, Host Conversations, Conversations, Response Time Details Views.

The From and To fields are enabled only when the Time Range is set to Custom.

Saving Filter Parameters

After clicking the Filter button in the Interactive Report and selecting the desired parameters, you can then save these selections with the purpose of viewing that same data at a future time. Enter a name in the Filter Name field. A filter is saved only be saved if a filter name is entered. Only saved filters are persisted across multiple login sessions. Click the **Submit** button.

This filter is now saved and displayed underneath the Interactive Report. You can save up to five filters.

Setting up Scheduled Exports

You can create a Scheduled Export to have the dashboards extracted regularly and sent to you in CSV or HTML format.

You can set up scheduled jobs that will generate a daily report at a specified time, in the specified interval, and then e-mail it to a specified e-mail address. You can also obtain a report on the spot clicking on the Preview button, rather than wait for the scheduled time. This report can also be sent after you preview it.

To set up a Scheduled Export, follow these steps:

Step 1 On most screens under **Network Analysis**, the Interactive Report is available on the left side of the screen. Click the **Export** button in the **Interactive Report** box. The **Create Scheduled Report** window appears.

Step 2 Choose the Export Type (Daily or Weekly).

Step 3 Choose the Export Time (when you would like the report delivered to you): Day and Hour.

Step 4 Choose the Report Time (if Daily) or the Data Time Range (if Weekly). This is the interval of time you would like measured.

Step 5 The Report Time for a daily report is restricted to the current 24 hours.

Step 6 The Report Time for a weekly report is always from 17:00 to 17:00, for however many days chosen.

For example:

If you choose Export Type “Weekly,” Data Time Range “Last 2 Days,” and Export Time: Day “Wednesday” and Hour “13:00,” the report will show data from Sunday at 17:00 to Tuesday at 17:00.

If you choose Export Time: Day “Wednesday” and Hour “18:00,” the report will show data from Monday at 17:00 to Wednesday at 17:00.

Step 7 Enter the e-mail address to which you would like the report delivered.

Step 8 Choose the delivery option (HTML or CSV).

Step 9 Enter the report description, which will appear at the end of the filename of the report delivered to you.

Step 10 Click:

- The **Reset** button to clear the values in the dialog box
- The **Preview** button to preview the report
- The **Submit** button to submit the request for the scheduled job
- The **Cancel** button to close the dialog box and return to the previous screen

Top Talkers Dashboards

This section includes the following topics:

Traffic Summary

The Top Talkers Summary dashboard allows you to view the Top N Applications, Top N Application Groups, Top N Hosts (In and Out), IP Distribution by Bytes, Top N DSCP, and Top N VLAN that is being monitored on your network. It provides auto-monitoring of traffic from all WAAS devices. You can view the Traffic Summary Dashboard by choosing **Monitor > Network Analysis Module > Overview.**

You can use the Interactive Report on the left to filter the information for a particular site, data source, VLAN, or reporting time interval. You can specify just one type of criteria and leave the others blank, or specify all of them. You can also choose to view the rate or cumulative data from the Interactive Report.

When you log into the NAM for the first time, the default view is the Traffic Summary dashboard, and the top data source is selected by default.

The charts shown on this dashboard are as follows:

- Top N Applications

  The Top N Applications Chart enables you to view the traffic rate (bytes per second or bits per second) or traffic volume (bytes or bits), depending on the Interactive Report filter selection (data rate or cumulative, respectively). When you place your cursor over the colored bar, you will see the number of bytes per second collected or the total bytes over the last time interval.

- Top N Application Groups

  This chart shows a detailed analysis of the Top N application groups and the traffic rate or volume for this interval. In the Interactive Report, you can select either rate or cumulative, where rate is the bytes per second, and cumulative is the total number of bytes.

- Top N Hosts (In and Out)

  This chart displays the traffic rate (bytes per second or bits per second) or traffic volume (bytes or bits).

- IP Distribution by Bytes

  This chart shows the percentages of bytes that are distributed to IP protocols (for example, IPv4 TCP).

- Top N DSCP

  This chart shows statistics for the top DSCP aggregation groups.

- Top N VLAN

  This chart shows the Top N VLAN statistics. In this chart, you might see VLAN 0, which is for traffic that does not have any VLAN tags.

To see a chart in table format, use the View as Chart / View as Grid toggle button on the bottom right corner of the chart. You can also click the **Show as Image** button to view the image and save it as a PNG file.

When viewing the data as a Grid, the numbers are formatted according to what you have configured in Configure > Network Analysis Module > Advanced > Preferences. In that window, you can also configure the number of Top N entries you would like to display.

Top Talkers Details

While you are in the process of deploying WAAS devices, you can get data to assist in the WAAS planning and configuration. For information about setting up WAN traffic, see Adding a Data Source for a New WAAS Device.

When you choose **Monitor > Network Analysis Module > Top Talkers Details**, you will see the window that assists you in the predeployment process. Use the Interactive Report window to select the traffic you want to analyze for optimization. The window displays the Top Applications, Top Network Links, Top Clients, and Top Servers.

Based on the results, you can then configure the WAAS products to optimize your network.

Throughput Dashboards

This section includes the following topics:

- Network, page 15-18
- Top Applications, page 15-18
- Application, page 15-19

Network

The Network dashboard enables you to view LAN versus WAN throughput for WAAS users both in the incoming and outgoing directions. To view these reports, configure interface groups that comprise WAN and LAN interfaces. The displayed information represents the total data collected since the collection was created, or since the NAM was restarted. To view the Network dashboard, choose **Monitor > Network Analysis Module > Throughput > Network**.

Choose an interface group view from the Interface Selector on the left side of the window to see traffic in the charts. Click the arrow icon to the left of the NDE data source name to display all interfaces groups, and then select an interface group view. If the charts show no data, and you see the message "Interface needs to be selected," you have not yet chosen an interface group view.

Once chose the interface group view, you see the following charts populated:

- Interface Traffic (Ingress % Utilization and Egress % Utilization)
- Top N Applications—Ingress
- Top N Applications—Egress
- Top N Hosts—Ingress
- Top N Hosts—Egress
- Top N DSCP Aggr—Ingress
- Top N DSCP Aggr—Egress

You can enter the interface speed manually through the Interface capacity table, or the speed can be auto configured if the SNMP settings for the NDE device are entered in the data source table.

Top Applications

In the Top Applications dashboard, you can view the top applications by the traffic rate over a selected time and for the specified site and/or data source.

Applications Over Time shows you all of the applications that have been running for the time period interval. The color-coded legend shows you what the applications are running.

If you place your cursor over any of the data points, you get more details about the exact values for each of the applications that are running.

Application

- In the Application window, you can see the traffic level for a given application over a selected period of time. It is available under the **Monitor > Network Analysis Module > Throughput > Application**. This window shows you the following:
- A graph of application traffic over time.
- Top hosts that transmit and receive traffic on that application for the selected time period.
- Application Configuration that shows the criteria by which the NAM classifies packets as that application. This criteria is typically a list of TCP and/or UDP ports that identify the application. Note that some applications are identified by heuristic or other state-based algorithms.

Performance Analysis Dashboards

This section includes the following topics:

Application

The Application dashboard provides the transaction time performance for an application as well as the original and optimized traffic volume reported by the flow agent. Information about how the transaction time is broken up across client, WAN, and server segments is also provided. For example, if the transaction time is dominated by the server segment time (due to a slow server), WAAS may not be able to improve the performance as much as when it is dominated by WAN network time. To view the Application performance analysis dashboard, choose **Monitor > Network Analysis Module > Performance Analysis > Application**.

The charts available on this dashboard are as follows:

- Transaction Time (Client Experience)
- Traffic Volume and Compression Ratio
- Average Concurrent Connections (Optimized vs. Passthru)
- Multi-Segment Network Time (Client LAN - WAN - Server LAN)

Conversation Multisegments

The Conversation Multiple Segments dashboard correlates data from different data sources, and allows you to view and compare response time metrics from multiple WAAS segments (data sources). To view the Conversation Multiple Segments dashboard, choose **Monitor > Network Analysis Module > Performance Analysis > Conversation Multisegments**.

The Response Time Across Multiple Segments window shows the response time metrics of the selected server or client-server pair from applicable data sources.

CHAPTER 16

Maintaining Your WAAS System

This chapter describes the tasks that you may need to perform to maintain your WAAS system.

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

This chapter contains the following sections:

Upgrading the WAAS Software

Table 16-1 outlines the steps you must complete to upgrade your WAAS software to a more recent version.

We recommend that all devices in your WAAS network should be running the same version of the WAAS software. If some of your WAAS devices are running different software versions, the WAAS Central Manager should be the highest version. For details on version interoperability limitations, see the *Release Note for Cisco Wide Area Application Services*.

If the WAAS Central Manager sees any registered WAE devices that are at a higher version level, it raises a minor alarm to alert you. Additionally, the WAE devices are shown in red on the device listing page.

WAAS Central Manager version 5.0.1 can manage WAE devices that are running version 4.2.1 and later releases. Some WAAS Central Manager windows (with new features) are not applicable to WAAS devices that are running a version lower than 5.0.1. If you modify the configuration in such windows, the configuration is saved, but it has no effect on the device until the device is upgraded to version 5.0.1.

Note WAAS version 5.0 is not supported running in a mixed version WAAS network where any WAAS device is running a software version lower than 4.2.1. If you have any WAAS devices running versions earlier than 4.2.1, you must first upgrade them to version 4.2.1 (or a later version) before you install version 5.0 on the Central Manager. Do not upgrade any device to a version later than the existing Central Manager version. After all devices and the Central Manager are running version 4.2.1 or later, then you can begin the upgrade to version 5.0 on the WAAS Central Manager. Directly upgrading a device from version 4.0 or 4.1 to 5.0 is not supported.

Upgrading is supported only from certain older releases to a particular release. If you have a WAAS device that is running a release from which upgrading to the desired release is not supported, first upgrade the device to an intermediate supported release and then to the final desired release. For details on what versions are supported for upgrades, see the *Release Note for Cisco Wide Area Application Services* for the software version to which you want to upgrade.

Table 16-1 Checklist for Upgrading the WAAS Software

Task	Additional Information and Instructions
1. Determine the current software version running on your WAAS network.	Check the software version that you are currently using so when you go to Cisco.com, you know if there is a newer version to download. For more information, see the "Determining the Current Software Version" section on page 16-3.
2. Obtain the new WAAS software version from Cisco.com.	Visit Cisco.com to download a newer software version and place this file on a local FTP or HTTP server. For more information, see the "Obtaining the Latest Software Version from Cisco.com" section on page 16-3.
3. Register the new software version with the WAAS Central Manager.	Register the URL of the new software file so the WAAS Central Manager knows where to go to access the file. For more information, see the "Specifying the Location of the Software File in the WAAS Central Manager GUI" section on page 16-3.
4. Upgrade your WAAS Central Manager.	Upgrade the standby and primary WAAS Central Managers. For more information, see the "Upgrading the WAAS Central Manager" section on page 16-5.
5. Upgrade your WAAS devices using Device Groups.	After upgrading the WAAS Central Manager, upgrade all your WAAS devices that are members of a device group. For more information, see the "Upgrading Multiple Devices Using Device Groups" section on page 16-7.
6. Delete the software version file.	After completely upgrading your WAAS network, you can remove the software file if desired. For more information, see the "Deleting a Software File" section on page 16-8.

If you need to downgrade or roll back the WAAS software to a lower version, first downgrade or roll back the WAE devices, then the standby Central Manager (if applicable), and finally the primary Central Manager. For more information about downgrading, see the *Release Note for Cisco Wide Area Application Services* for your software version.

Determining the Current Software Version

To view the current software version running on any particular device, choose **Devices > All Devices**. The All Devices window displays the software version for each device listed.

You can also click **Devices >** *device-name* or the **Edit** icon next to the name of a device in the Devices window. The Device Dashboard window appears, listing the software version for that device.

Note The software version is not upgraded until a software upgrade has been successfully completed. If a software upgrade is in progress, the version number displayed is the base version, not the upgraded version number.

Alternatively, in the device context, choose **Monitor > CLI Commands > Show Commands**. Choose **version** and click **Submit**. A secondary window pops up and displays the CLI output for the **show version** command.

Obtaining the Latest Software Version from Cisco.com

To obtain the latest WAAS software version from Cisco.com, follow these steps:

Step 1 Launch your web browser and open this location:

http://www.cisco.com/cisco/software/navigator.html

Step 2 Navigate to the **Application Networking Services > Wide Area Application Services > Cisco Wide Area Application Services (WAAS) Software** download area.

Step 3 Choose the WAAS software version that you want and download the appropriate software image (Universal, Accelerator only, or No Payload Encryption [NPE]).

Step 4 Register the location of the software file in the WAAS Central Manager GUI, as described in the section that follows.

Specifying the Location of the Software File in the WAAS Central Manager GUI

To upgrade your WAAS software, you must first specify the location of the WAAS software file in the WAAS Central Manager GUI and configure the software file settings.

There are two types of WAAS software files, as follows:

- Universal—Includes Central Manager, Application Accelerator, and AppNav Controller functionality. You can use this type of software file to upgrade a device operating in any mode.
- Accelerator only—Includes Application Accelerator and AppNav Controller functionality only. You can use this type of software file to upgrade only an Application Accelerator or AppNav Controller device. If you want to change an Application Accelerator or AppNav Controller to a Central

Manager, you must install the Universal software file, reload the device, change the device mode to central-manager, and then reload the device again. Additionally, kdump analysis functionality is not included in the Accelerator only image.

To configure the software file settings form, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > Version Management > Software Update**.

Step 2 Click the **Create New Software File** icon in the taskbar.

The Creating New Software File window appears. (See Figure 16-1.)

Figure 16-1 Creating New Software File Window

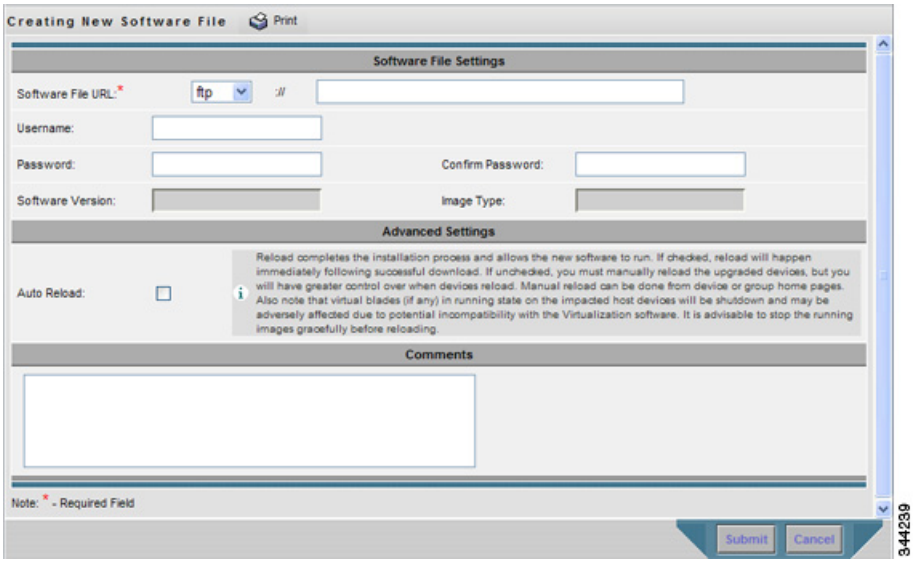

Step 3 In the Software File URL field, specify the location of the new WAAS software file as follows:

a. Choose a protocol (**http** or **ftp**) from the drop-down list.

b. Enter the URL for the .bin software file that you downloaded from Cisco.com. For example, a valid URL might look like the following:

http://internal.mysite.com/waas/*WAAS-xxxx-K9*.bin

where WAAS-*xxxx*-K9.bin is the name of the software upgrade file. (The filename typically includes the version number.)

Be sure that the URL identifies the correct type of software image for the devices you want to upgrade, either Universal or Accelerator only.

Step 4 If your server requires user login authentication, enter your username in the Username field and enter your login password in the Password field. Enter the same password in the Confirm Password field.

The Software Version and Image Type fields cannot be edited. They are filled in automatically after you submit the settings and the image is validated.

Step 5 In the Advanced Settings section, check the **Auto Reload** check box to automatically reload a device when you upgrade the software. If you do not check this box, you will need to manually reload a device after you upgrade the software on it to complete the upgrade process.

Note During a device reload, any virtual blades running on the device are shut down and may be adversely affected due to a potential incompatibility with the virtualization software. Therefore, you should stop the running images gracefully before reloading.

Step 6 (Optional) Enter comments in the field provided.

Step 7 Click **Submit**.

The software image file is validated and the Software Version and Image Type fields are filled in with the appropriate information extracted from the image file.

Caution If your browser is configured to save the username and password for the WAAS Central Manager GUI, the browser will autopopulate the username and password fields in the Creating New Software File window. You must clear these fields before you click **Submit**.

The software file that you want to use is now registered with the WAAS Central Manager. When you perform the software upgrade or downgrade, the URL that you just registered becomes one of the choices available in the Update Software window.

To reload a device from the CLI, use the **reload** EXEC command.

Note When you are viewing the list of registered software files, if the Image Type column shows Unknown for a software file, it indicates that the software file was added under a WAAS version previous to 4.2.1. These Unknown software files must be resubmitted if you want to use them. Click the **Edit** icon next to the file to open the Modifying Software File window, and then click the **Submit** button to resubmit the file.

Upgrading the WAAS Central Manager

When upgrading software in your WAAS network, begin with WAAS Central Manager before upgrading the WAE devices.

Primary and standby WAAS Central Manager devices must be running the same version of WAAS software. If they are not, the standby WAAS Central Manager detects this and will not process any configuration updates it receives from the primary WAAS Central Manager. If the primary WAAS Central Manager sees that the standby WAAS Central Manager has a different version level, it shows the standby WAAS Central Manager in red on the device listing page.

If you use the primary WAAS Central Manager to perform the software upgrade, you need to upgrade your standby WAAS Central Manager first, and then upgrade your primary WAAS Central Manager. We also recommend that you create a database backup for the primary WAAS Central Manager and copy the database backup file to a safe place before you upgrade the software.

Use this upgrade procedure for WAAS Central Manager devices. You can also use this upgrade procedure to upgrade WAAS devices one at a time (after the WAAS Central Manager).

To upgrade your software to another WAAS software release on a single device, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

The Device Dashboard window appears.

Step 2 Verify that the device is not already running the version to which you plan to upgrade.

Step 3 Click the **Update** link.

The Software Update window appears.

Step 4 Choose the software file URL from the Software Files list by clicking the radio button next to the filename.

The list displays only software files with an image type of Universal, because you are upgrading a Central Manager device. If no such images are available, you must create a software file as described in the "Specifying the Location of the Software File in the WAAS Central Manager GUI" section on page 16-3.

Step 5 Click **Submit**, and then click **OK** to confirm your decision.

The Devices listing window reappears. You can monitor the progress of your upgrade from this window.

Software upgrade status messages are displayed in the Software Version column. These intermediate messages are also written to the system log on the WAAS devices. See Table 16-2 for a description of upgrade status messages.

Step 6 Clear your browser cache, close the browser, and restart the browser session to the WAAS Central Manager.

The WAAS Central Manager may reboot at the conclusion of the upgrade procedure (if Auto Reload was checked in the Creating New Software File window), causing you to temporarily lose contact with the device and the graphical user interface.

Table 16-2 Upgrade Status Messages

Upgrade Status Message	Condition
Pending	The request has yet to be sent from the WAAS Central Manager to the device, or receipt of the request has yet to be acknowledged by the device.
Downloading	The download method for the software file is being determined.
Proceeding with Download	The download method for the software file is determined to be direct download. Proceeding with the request for direct download of the software file.
Download in Progress (Completed …)	The direct download of the software file is being processed. "Completed" indicates the number of megabytes processed.
Download Successful	The direct download of the software file has been successful.
Download Failed	The direct download of the software file cannot be processed. Further troubleshooting is required; see the device system message log. If you are upgrading several devices at once, the download may fail if the server hosting the software file becomes overloaded with requests. Retry the upgrade by clicking the Retry link if shown.
Proceeding with Flash Write	A request has been made to write the software file to the device flash memory.

Table 16-2 Upgrade Status Messages (continued)

Upgrade Status Message	Condition
Flash Write in Progress (Completed …)	The write of the device flash memory is being processed. "Completed" indicates the number of megabytes processed.
Flash Write Successful	The flash write of the software file has been successful.
Reloading	A request to reload the device has been made in order to complete the software upgrade. The device may be offline for several minutes.
Reload Needed	A request to reload the device has not been made. The device must be reloaded manually to complete the software upgrade.
Cancelled	The software upgrade request was interrupted, or a previous software upgrade request was bypassed from the CLI.
Update Failed	The software upgrade could not be completed. Troubleshooting is required; see the device system message log. If you are upgrading several devices at once, the upgrade may fail if the server hosting the software file becomes overloaded with requests. Retry the upgrade by clicking the Retry link if shown.

Upgrading Multiple Devices Using Device Groups

Note This procedure is for WAE devices only. WAAS Central Manager devices cannot be upgraded using device groups.

To upgrade to a more recent WAAS software release on multiple devices, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Device Groups >** *device-group-name*.

Step 2 Choose **Admin > Versioning > Software Update**.

The Software Update for Device Group window appears.

Step 3 Choose the software file URL from the Software File URL list by clicking the radio button next to the filename. If no images are available, you must create a software file as described in the "Specifying the Location of the Software File in the WAAS Central Manager GUI" section on page 16-3.

If you are updating many devices and you want to use a smaller size software file to save network bandwidth, specify a software file with an image type of Accelerator only, which is smaller than a Universal image. If you later want to change an Accelerator-only device to a Central Manager, you must install the Universal software file, reload the device, change the device mode to central-manager, and then reload the device again.

Step 4 Click **Submit**.

To view the progress of an upgrade, go to the All Devices window (**Devices > All Devices**) and view the software upgrade status message in the Software Version column. These intermediate messages are also written to the system log on WAAS devices. See Table 16-2 for a description of the upgrade status messages.

Deleting a Software File

After you have successfully upgraded your WAAS devices, you can remove the software file from your WAAS system.

Note You may want to wait a few days before removing a software file in the event you need to downgrade your system for any reason.

To delete a WAAS software file, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > Version Management > Software Update**.

Step 2 Click the **Edit** icon next to the software file that you want to delete. The Modifying Software File window appears.

Step 3 Click the **Trash** icon in the taskbar.

You are prompted to confirm your decision to delete the software file.

Step 4 Click **OK**.

You are returned to the Software Files listing window with the selected software file removed from the WAAS network.

Backing Up and Restoring your WAAS System

This section contains the following topics:

Backing Up and Restoring the WAAS Central Manager Database

The WAAS Central Manager device stores WAAS network-wide device configuration information in its Centralized Management System (CMS) database. You can manually back up the CMS database contents for greater system reliability.

The CMS database backup is in a proprietary format that contains an archive database dump, WAAS Central Manager registration information, and device information that the WAAS Central Manager uses to communicate with other WAAS devices. CMS database backup files are not interchangeable between primary and standby WAAS Central Manager devices. This means you cannot use the backup file from a primary WAAS Central Manager to restore a standby WAAS Central Manager.

To back up the CMS database for the WAAS Central Manager, use the **cms database backup** EXEC command. For database backups, you need to specify the location, password, and user ID of the remote server that you want to store the backup file.

Note If you have a backup made when the secure store was in user-passphrase mode and you restore it to a system where the secure store is in auto-passphrase mode, you must enter the user passphrase to proceed with the restore. After the restore, the system is in user-passphrase mode. If you have a backup made when the secure store was in auto-passphrase mode and you restore it to a system where the secure store is in user-passphrase mode, you do not need to enter a password. After the restore, the system is in auto-passphrase mode.

To back up and restore the CMS database, follow these steps:

Step 1 On the WAAS Central Manager GUI device, use the **cms database backup** command to back up the CMS database to a file, as shown in the following example:

```
CM# cms database backup
Creating database backup file cms-db-11-05-2010-15-22_4.3.1.0.1.dump
Backup file cms-db-11-05-2010-15-22_4.3.1.0.1 is ready.
Please use `copy' commands to move the backup file to a remote host.
```

Note The backup file is automatically given a name in the following format cms-db-*date-timestamp_version*.dump. For example, cms-db-7-22-2010-17-36_4.3.1.0.1.dump. Note that the timestamp is in 24-hour format (HH:MM) that does not show seconds.

Step 2 Save the file to a remote server by using the **copy disk ftp** command.

This command copies the file from the local disk to a remote FTP server, as shown in the following example:

```
CM# cd /local1
CM# copy disk ftp 10.86.32.82 /incoming cms-db-7-22-2008-17-36_4.1.3.0.1.dump
cms-db-7-22-2008-17-36_4.1.3.0.1.dump

Enter username for remote ftp server:ftp
Enter password for remote ftp server:*******
Initiating FTP upload...
Sending:USER ftp
10.86.32.82 FTP server (Version wu-2.6.1-18) ready.
Password required for ftp.
Sending:PASS *******
User ftp logged in.
Sending:TYPE I
```

```
Type set to I.
Sending:PASV
Entering Passive Mode (10,86,32,82,112,221)
Sending:CWD /incoming
CWD command successful.
Sending PASV
Entering Passive Mode (10,86,32,82,203,135)
Sending:STOR cms-db-7-22-2008-17-36_4.1.3.0.1.dump
Opening BINARY mode data connection for cms-db-7-22-2008-17-36_4.1.3.0.1.dump.
Transfer complete.
Sent 18155 bytes
```

Step 3 Restore the CMS database as follows:

a. Disable the CMS service:

```
CM# configure
CM(config)# no cms enable
CM(config)# exit
```

Note Stopping the CMS service disables the WAAS Central Manager GUI. All users currently logged into this GUI are automatically logged out once the CMS service is disabled.

b. Delete the existing CMS database:

```
CM# cms database delete
```

c. Initialize the CMS database:

```
CM# cms database create
```

d. Restore the CMS database contents from the backup file:

```
CM# cms database restore cms-db-7-22-2008-17-36_4.1.3.0.1.dump
```

Note After the restore, any WAEs that were registered with the Central Manager during the time since the backup was created will be disconnected from the Central Manager because there is no information about them in the backup file. To bring these WAEs online, you must deregister and reregister them with the Central Manager. On each WAE that was disconnected, use the following commands:

```
WAE# cms deregister force
WAE# configure
WAE(config)# cms enable
```

e. Enable the CMS service on the Central Manager:

```
CM# configure
CM(config)# cms enable
```

Backing Up and Restoring a WAE Device

You should back up the database of each WAAS device on a regular basis in case a system failure should occur.

Note The backup and restore methods described in this section apply only to a WAE device that is not configured as a WAAS Central Manager. For information on backing up the WAAS Central Manager device, see the "Backing Up and Restoring the WAAS Central Manager Database" section on page 16-9.

You can use either of the following methods to back up and restore the database of an individual WAE device:

- WAE Device Manager—For information on using the WAE Device Manager to back up and restore a device's database, see the "Backing Up the Configuration Files" section on page 11-6.
- CLI—You can use the **copy running-config** command to back up and restore a device's configuration. This command saves the currently running configuration.

Additionally, you can restore a WAE to the default configuration that it was manufactured with at any time by removing the user data from the disk and Flash memory, and erasing all existing files cached on the appliance. Basic configuration information, such as network settings, can be preserved. The appliance is accessible through Telnet and Secure Shell (SSH) after it reboots.

Note If software upgrades have been applied, the restoration process returns to the defaults of the currently installed version and not the factory defaults.

To restore a WAE to its factory defaults or the defaults of the current configuration from the CLI, use the **restore factory-default** [**preserve basic-config**] EXEC command.

For more information about the CLI commands, see the *Cisco Wide Area Application Services Command Reference*.

Reinstalling the System Software

This section contains instructions for using the software recovery files to reinstall your system software if for some reason the software that is installed has failed. A software recovery CD-ROM ships with some WAE and WAVE hardware devices and some WAVE devices use a USB flash drive for recovery.

Caution If you upgraded your software after you received your software recovery CD-ROM or image files, using the recovery software images may downgrade your system. Ensure that you are using the desired software recovery version.

The WAAS software consists of three basic components:

- Disk-based software
- Flash-based software
- Hardware platform cookie (stored in flash memory)

All of these components must be correctly installed for WAAS software to work properly.

The software is contained in two types of software images provided by Cisco Systems:

- A .bin image that contains disk and flash memory components (the Universal version of the WAAS software)
- A .sysimg image that contains a flash memory component only

An installation that contains only the WAAS flash memory-based software, without the corresponding disk-based software, boots and operates in a limited mode, allowing for further disk configuration before completing a full installation.

The .sysimg component is provided for recovery purposes and allows for repair of flash memory only without modifying the disk contents.

Note The system image used depends on your device. For all WAVE and WAE-674/7341/7371 devices (64-bit platforms), use the 64-bit system image (with "x86_64" in its name). For all other devices, use the 32-bit system image named without this designator.

An NPE image is provided that has the disk encryption feature disabled for use in countries where disk encryption is not permitted.

If you have a WAVE appliance that requires a USB flash drive for software recovery, your USB flash drive must contain both of the needed software images in the form of an ISO archive file that you copy to the flash drive. (See the "Preparing the USB Flash Drive" section on page 16-13).

These options are available from the software recovery installer menu:

- Option 1: Configure Network—If the .bin image you need to install is located on the network instead of the CD-ROM or USB flash drive (which may be the case when an older CD-ROM or USB image is used to install new software), then you must choose this option to configure the network before attempting to install the .bin image.

  This option is performed automatically if you install a .sysimg file from the network.
- Option 2: Manufacture Flash—This option verifies the flash memory and, if invalid, automatically reformats it to contain a Cisco standard layout. If reformatting is required, a new cookie is installed automatically.

  This option is performed automatically as part of a .bin or .sysimg installation.
- Option 3: Install Flash Cookie—This option generates a hardware-specific platform cookie and installs it in flash memory. You need to use this option only if there has been a change in the hardware components, such as replacing the motherboard, or if you moved a flash memory card between systems.

  This option is performed automatically during the flash manufacturing process, if needed, as part of a .bin or .sysimg installation.
- Option 4: Install Flash Image from Network and Option 5: Install Flash Image from USB/CD-ROM —These options allow installation of the flash memory .sysimg only and do not modify disk contents. They may be used when a new chassis has been provided and populated with the customer's old disks that need to be preserved.

  These options automatically perform flash verification and hardware cookie installation, if required. When installing from the network, you are prompted to configure the network if you have not already done so.
- Option 6: Install Flash Image from Disk—This option is reserved for future expansion and is not available.
- Option 7: Recreate RAID device—This option applies only to WAE-674, WAE-7341, WAE-7371, WAVE-7541, WAVE-7571, and WAVE-8541 devices and recreates the RAID array.
- Option 8: Wipe Out Disks and Install .bin Image—This option provides the preferred procedure for installing the WAAS software.

Caution Option 8 erases the content from all disk drives in your device.

This option performs the following steps:

a. Checks that flash memory is formatted to Cisco specifications. If yes, the system continues to step b. If no, the system reformats the flash memory, which installs the Cisco file system, and generates and installs a platform-specific cookie for the hardware.

b. Erases data from all drives.

c. Remanufactures the default Cisco file system layout on the disk.

d. Installs the flash memory component from the .bin image.

e. Installs the disk component from the .bin image.

Note When option 8 is used and the system reboots and begins optimizing traffic, the show **disks details command** may show that the /dre1 partition is 98% or more used, due to the preallocation of DRE cache space. Use the **show statistics dre** command to display the actual DRE cache usage.

- Option 9: Exit (reboot)—This option ejects the CD-ROM (if applicable) and reboots the device. If you are using a USB flash drive for software installation, remove it from the device before rebooting.

The following sections describe how to reinstall the WAAS system software:

- "Preparing the USB Flash Drive" section on page 16-13—Read this section if you have a WAVE appliance that requires a USB flash drive instead of a CD to install the system software.
- "Reinstalling the System Software" section on page 16-15—Describes how to reinstall the system software from a CD or USB flash drive.
- "Ensuring RAID Pairs Rebuild Successfully" section on page 16-18—Describes how to ensure that RAID disks rebuild successfully.

Preparing the USB Flash Drive

If you have a WAVE appliance that requires a USB flash drive for software recovery, you must prepare the USB flash drive with the appropriate files before you can start the software recovery process. You will need the following:

- Windows PC (Windows XP or 7) or Mac computer
- USB flash drive that is 1 GB or larger in size
- The following software recovery files:
  - WAAS Rescue CD ISO image file, which is available in the WAAS Software Download area of Cisco.com. The filename is similar to waas-rescue-cdrom-*x.x.x.x*-k9.iso, where the *x*'s denote the software version number. Alternatively, the ISO image file is available on the WAAS release DVD, or you can make an ISO image file from a WAAS recovery CD.
  - syslinux.cfg file, which is also available in the WAAS Software Download area of Cisco.com and on the WAAS release DVD.
  - Unetbootin utility for Windows, which is available from the Unetbootin Sourceforge website. (This file is not needed if you are using a Mac computer to prepare the USB flash drive.)

The steps for preparing the USB flash drive differ if you are using a Windows or Mac computer. See the appropriate topic:

- "Using a Windows Computer to Prepare the USB Flash Drive" section on page 16-14
- "Using a Mac Computer to Prepare the USB Flash Drive" section on page 16-14

Using a Windows Computer to Prepare the USB Flash Drive

To prepare the USB flash drive on a Windows PC, follow these steps:

Step 1 Transfer the software recovery files onto the Windows computer, noting the directory where they are stored.

Step 2 Insert the USB flash drive into a USB port on the PC.

Step 3 Open My Computer.

Step 4 Right click on the Removable Disk (drive letter will vary with system) and select **Format**.

In the formatting tool, select FAT32 for the File System and check the Quick Format check box, and then click **Start**. Click OK on the warning message. Close the formatting tool after the formatting is complete.

Step 5 Double-click the Windows Unetbootin utility to run it.

Step 6 Select the Diskimage option and click the corresponding browse button (...) to select the waas-rescue-cdrom-*x.x.x.x*-k9.iso image file.

Step 7 Ensure that USB Drive is selected in the Type pull-down list and that the correct drive letter is selected for Drive.

Step 8 Click OK to install the bootable image on the USB flash drive. When the installation has completed, click **Exit**.

Step 9 Drag copy the syslinux.cfg file onto the USB flash drive and confirm to replace. This file replaces the existing file on the USB flash drive with one customized for your WAAS system.

Step 10 Remove the USB flash drive from the PC.

To continue reinstalling the system software from the prepared USB flash drive, follow the instructions in the "Reinstalling the System Software" section on page 16-15.

Using a Mac Computer to Prepare the USB Flash Drive

Step 1 Transfer the software recovery files onto the Mac computer, noting the directory where they are stored.

Step 2 Open the Terminal application.

Step 3 Enter the **diskutil** command to list the current storage devices.

Step 4 Insert the USB flash drive into a USB port on the Mac computer.

Step 5 Enter the **diskutil** command again to list the current storage devices and note the device node assigned to your USB flash drive (for example, /dev/disk2)

Step 6 Enter the **diskutil unmountDisk /dev/disk*N*** command to unmount the USB flash drive. (*N* denotes the number of the USB flash drive node.)

Step 7 Enter the command **sudo dd if=/*path*/waas-rescue-cdrom-*x.x.x.x*-k9.iso of=/dev/disk*N* bs=1m** to install the bootable image on the USB flash drive. (*path* denotes the folder path to the WAAS ISO file, *x.x.x.x* denotes the WAAS software version number, and *N* denotes the number of the USB flash drive node.)

Step 8 You may receive a warning message about the sudo command and a prompt to enter your password to proceed. Enter your password if required.

Step 9 Copy the syslinux.cfg file onto the USB flash drive.

Step 10 Enter the **diskutil eject /dev/disk*N*** command and remove your USB flash drive when the command completes.

To continue reinstalling the system software from the prepared USB flash drive, follow the instructions in the next section, Reinstalling the System Software.

Reinstalling the System Software

To reinstall the system software on a WAE appliance using the software recovery CD-ROM or USB flash drive, follow these steps:

Step 1 Connect a serial console to the WAAS appliance and use the console for the following steps.

Step 2 Insert the software recovery CD-ROM in the CD drive of the WAE device or, if the device uses a USB flash drive for recovery, insert a bootable USB flash drive with the software recovery files into the USB port of the device (see the "Preparing the USB Flash Drive" section on page 16-13). WAVE-294/594/694/7541/7571/8541 devices do not have CD drives and instead use a USB flash drive for software recovery.

Step 3 Reboot the WAE. During the boot process, the boot loader pauses for 30 seconds and you must choose the VGA console if you are using vWAAS. The prompt is shown as follows:

```
Type "serial" for WAE/WAVE appliance.
Type "vga" for vWAAS.
boot:
```

Enter the **vga** command at the prompt to continue the boot process for the VGA console on vWAAS. After 30 seconds with no input, the boot process continues with the standard serial console for WAAS appliances.

After the WAE boots, you see the following menu:

```
Installer Main Menu:
  1. Configure Network
  2. Manufacture flash
  3. Install flash cookie
  4. Install flash image from network
  5. Install flash image from usb/cdrom
  6. Install flash image from disk
  7. Recreate RAID device (WAE-674/7341/7371/7541/7571/8541 only)
  8. Wipe out disks and install .bin image
  9. Exit (reboot)
Choice [0]:
```

Note The option numbers in the installer main menu may vary, depending on the WAAS software release being installed.

Step 4 Choose option 2 to prepare the flash memory.

This step prepares a cookie for the device and also retrieves the network configuration that was being used by the WAAS software. This network configuration is stored in the flash memory and is used to configure the network when the WAAS software boots up after installation.

Step 5 Choose option 3 to install the flash cookie that you prepared in the previous step.

Step 6 Choose option 5 to install the flash image from a CD-ROM or USB flash drive.

Step 7 (Optional) If you are working with a WAE-674, WAE-7341, WAE-7371, WAVE-7541, WAVE-7571, or WAVE-8541 device, choose option 7 to recreate the RAID array.

Step 8 Choose option 8 to wipe the disks and install the binary image.

This step prepares the disks by erasing them. The WAAS software image is installed.

Step 9 If you are using a USB flash drive to install the software, remove it from the device.

Step 10 Choose option 9 to reboot the WAE.

After the WAE reboots, it is running the newly installed WAAS software. The WAE has a minimal network configuration and is accessible via the terminal console for further configuration.

To reinstall the system software on an NME-WAE network module installed in a Cisco access router, follow these steps:

Step 1 Log in to the Cisco router in which the NME-WAE module is installed, and reload the NME-WAE module:

```
router-2851> enable
router-2851# service-module integrated-Service-Engine 1/0 reload
```

Step 2 Immediately open a session on the module:

```
router-2851# service-module integrated-Service-Engine 1/0 session
```

Step 3 While the module is reloading, you will see the following option during boot phase 3. Enter *** as instructed:

```
[BOOT-PHASE3]: enter `***' for rescue image: ***
```

Step 4 The rescue image dialog appears. The following example demonstrates how to interact with the rescue dialog (user input is denoted by entries in bold typeface):

```
This is the rescue image.  The purpose of this software is to let
you install a new system image onto your system's boot flash
device.  This software has been invoked either manually
(if you entered `***' to the bootloader prompt) or has been
invoked by the bootloader if it discovered that your system image
in flash had been corrupted.

To download an image from network, this software will request
the following information from you:
    - which network interface to use
    - IP address and netmask for the selected interface
    - default gateway IP address
    - FTP server IP address
    - username and password on FTP server
    - path to system image on server

Please enter an interface from the following list:
   0: GigabitEthernet 1/0
```

```
   1: GigabitEthernet 2/0
enter choice: 0
Using interface GigabitEthernet 1/0

Please enter the local IP address to use for this interface:
[Enter IP Address]: 10.1.13.2

Please enter the netmask for this interface:
[Enter Netmask]: 255.255.255.240

Please enter the IP address for the default gateway:
[Enter Gateway IP Address]: 10.1.13.1

Please enter the IP address for the FTP server where you wish
to obtain the new system image:
[Enter Server IP Address]: 10.107.193.240

Please enter your username on the FTP server (or 'anonymous'):
[Enter Username on server (e.g. anonymous)]: username

Please enter the password for username 'username' on FTP server:

Please enter the directory containing the image file on the FTP server:
[Enter Directory on server (e.g. /)]: /

Please enter the file name of the system image file on the FTP server:
[Enter Filename on server]: WAAS-5.0.1.10-K9.sysimg


Here is the configuration you have entered:

Current config:
          IP Address: 10.1.13.2
             Netmask: 255.255.255.240
     Gateway Address: 10.1.13.1
      Server Address: 10.107.193.240
            Username: username
            Password: *********
     Image directory: /
      Image filename: WAAS-5.0.1.10-K9.sysimg

Attempting download...
Downloaded 15821824 byte image file
A new system image has been downloaded.
You should write it to flash at this time.
Please enter 'yes' below to indicate that this is what you want to do:
[Enter confirmation ('yes' or 'no')]: yes
Ok, writing new image to flash
.............................................................................. done.
Finished writing image to flash.
Enter 'reboot' to reboot, or 'again' to download and install a new image:
[Enter reboot confirmation ('reboot' or 'again')]: reboot
Restarting system.
```

Step 5 After the module reboots, install the .bin image from an HTTP server:

```
NM-WAE-1# copy http install 10.77.156.3 /waas WAAS-5.0.1.10-k9.bin
```

Step 6 Reload the module:

```
NM-WAE-1# reload
```

After the module reboots, it is running the newly installed WAAS software.

Ensuring RAID Pairs Rebuild Successfully

You must ensure that all RAID pairs are done rebuilding before you reboot your WAE device. If you reboot while the device is rebuilding, you risk corrupting the file system.

RAID pairs will rebuild on the next reboot after you use the **restore factory-default** command, replace or add a hard disk drive, delete disk partitions, or reinstall WAAS from the booted recovery CD-ROM or USB flash drive.

To view the status of the drives and check if the RAID pairs are in "NORMAL OPERATION" or in "REBUILDING" status, use the **show disk details** EXEC command. When you see that RAID is rebuilding, you must let it complete that rebuild process. This rebuild process can take several hours.

If you do not wait for the RAID pairs to complete the rebuild process before you reboot the device, you may see the following symptoms indicating a problem:

- The device is offline in the Central Manager GUI.
- CMS can not be loaded.
- Error messages say that the file system is read-only.
- The syslog contains errors such as "Aborting journal on device md2," "Journal commit I/O error," "Journal has aborted," and "ext3_readdir: bad entry in directory."
- Other unusual behaviors related to disk operations or the inability to perform them.

If you encounter any of these symptoms, reboot the WAE device and wait until the RAID rebuild finishes normally.

Recovering the System Software

WAAS devices have a resident rescue system image that is invoked if the image in flash memory is corrupted. A corrupted system image can result from a power failure that occurs while a system image is being written to flash memory. The rescue image can help you download a system image to the main memory of the device and write it to flash memory.

Note The system image used depends on your device. For all WAVE and WAE-674/7341/7371 devices (64-bit platforms), use the 64-bit system image (with "x86_64" in its name). For all other devices, use the 32-bit system image named without this designator.

An NPE image is provided that has the disk encryption feature disabled for use in countries where disk encryption is not permitted.

To install a new system image using the rescue image, follow these steps:

Step 1 Download the system image file (*.sysimg) to a host that is running an FTP server.

Step 2 Establish a console connection to the WAAS device and open a terminal session.

Step 3 Reboot the device by toggling the power on/off switch.

After a few seconds, the bootloader pauses and prompts you to enter 1 to boot WAAS, r to boot the rescue image, x to reboot, or 9 to escape to the loader prompt. You have 10 seconds to respond before the normal boot process continues.

Step 4 Enter r to boot the rescue image.

The rescue image dialog appears and differs depending on whether your WAAS device was initially manufactured with version 4.x or 5.x. Step 5 describes the rescue image on a device that was initially manufactured with version 5.x. Step 6 describes the rescue image on a device that was initially manufactured with version 4.x.

Step 5 If you see the following output (from a device that was initially manufactured with version 5.x), login and use the **copy install** command to install the WAAS system software image (.bin file), as shown in the following example (user input is denoted by entries in bold typeface):

```
The device is running WAAS rescue image. WAAS functionality is unavailable
in a rescue image. If the rescue image was loaded by accident, please reload
the device. If the rescue image was loaded intentionally to reinstall WAAS software
please use the following command:

    copy [ftp|http|usb] install ...


SW up-to-date
...

Cisco Wide Area Virtualization Engine Console

Username: admin
Password:
System Initialization Finished.

WAVE# copy ftp install 172.16.10.10 / waas-universal-5.0.1.12-k9.bin
...

Installing system image to flash... Creating backup of database content before database
upgrade.
The new software will run after you reload.
WAVE# reload
Proceed with reload?[confirm]yes
Shutting down all services, will timeout in 15 minutes.
reload in progress ..Restarting system.
```

Step 6 If you see the following output (from a device that was initially manufactured with version 4.x), login and install the WAAS system image (.sysimg file), as shown in the following example (user input is denoted by entries in bold typeface):

```
This is the rescue image. The purpose of this software is to let
you download and install a new system image onto your system's
boot flash device. This software has been invoked either manually
(if you entered `***' to the bootloader prompt) or has been
invoked by the bootloader if it discovered that your system image
in flash had been corrupted.

To download an image, this software will request the following
information from you:
    - which network interface to use
    - IP address and netmask for the selected interface
    - default gateway IP address
    - server IP address
    - which protocol to use to connect to server
    - username/password (if applicable)
    - path to system image on server
```

```
Please enter an interface from the following list:
    0: GigabitEthernet 0/0
    1: GigabitEthernet 0/1
enter choice: 0
Using interface GigabitEthernet 0/0

Please enter the local IP address to use for this interface:
[Enter IP Address]: 172.16.22.22

Please enter the netmask for this interface:
[Enter Netmask]: 255.255.255.224

Please enter the IP address for the default gateway:
[Enter Gateway IP Address]: 172.16.22.1

Please enter the IP address for the FTP server where you wish
to obtain the new system image:
[Enter Server IP Address]: 172.16.10.10

Please enter your username on the FTP server (or 'anonymous'):
[Enter Username on server (e.g. anonymous)]: anonymous

Please enter the password for username 'anonymous' on FTP server:

Please enter the directory containing the image file on the FTP server:
[Enter Directory on server (e.g. /)]: /

Please enter the file name of the system image file on the FTP server:
[Enter Filename on server (e.g. WAAS-x86_64-4.x.x-K9.sysimg)]:
waas-x86_64-5.0.1.12-k9.sysimg

Here is the configuration you have entered:
Current config:
          IP Address: 172.16.22.22
             Netmask: 255.255.255.224
     Gateway Address: 172.16.22.1
      Server Address: 172.16.10.10
            Username: anonymous
            Password:
     Image directory: /
      Image filename: waas-x86_64-5.0.1.12-k9.sysimg

Attempting download...
Downloaded 31899648 byte image file
A new system image has been downloaded.
You should write it to flash at this time.
Please enter 'yes' below to indicate that this is what you want to do:
[Enter confirmation ('yes' or 'no')]: yes
Ok, writing new image to flash
Finished writing image to flash.
Enter 'reboot' to reboot, or 'again' to download and install a new image:
[Enter reboot confirmation ('reboot' or 'again')]: reboot
Restarting system.
Booting system, please wait.........
```

Step 7 Log in to the device as username **admin**. Verify that you are running the correct version by entering the **show version** command.

```
Username: admin
Password:

Console# show version
Cisco Wide Area Application Services Software (WAAS)
```

```
Copyright (c) 1999-2012 by Cisco Systems, Inc.
Cisco Wide Area Application Services (universal-k9) Software Release 5.0.1 (buil
d b12 May 28 2012)
Version: oe294-5.0.1.12

Compiled 23:23:45 May 28 2012 by damaster

Device Id: 50:3d:e5:9c:8f:a5
System was restarted on Tue May 29 16:35:50 2012.
System restart reason: called via cli.
The system has been up for 8 hours, 10 minutes, 19 seconds.
```

Recovering a Lost Administrator Password

If an administrator password is forgotten, lost, or misconfigured, you will need to reset the password on the device.

Note You cannot restore a lost administrator password. You must reset the password to a new one, as described in this procedure.

To reset the password, follow these steps:

Step 1 Establish a console connection to the device and open a terminal session.

Step 2 Reboot the device.

While the device is rebooting, watch for the following prompt, then press **Enter** when you see it:

```
Cisco WAAS boot:hit RETURN to set boot flags:0009
```

Step 3 When prompted to enter bootflags, enter the following value: **0x8000**

```
Available boot flags (enter the sum of the desired flags):
0x4000 - bypass nvram config
0x8000 - disable login security

[CE boot - enter bootflags]:0x8000
You have entered boot flags = 0x8000
Boot with these flags? [yes]:yes

[Display output omitted]
Setting the configuration flags to 0x8000 lets you into the system, bypassing all
security. Setting the configuration flags field to 0x4000 lets you bypass the NVRAM
configuration.
```

Step 4 When the device completes the boot sequence, you are prompted to enter the username to access the CLI. Enter the default administrator username (**admin**).

```
Cisco WAE Console

Username: admin
```

Step 5 When you see the CLI prompt, set the password for the user using the **username passwd** command in global configuration mode.

```
WAE# configure
WAE(config)# username admin passwd
```

This command invokes interactive password configuration. Follow the CLI prompts.

Step 6 Save the configuration change:

```
WAE(config)# exit
WAE# write memory
```

Step 7 (Optional) Reboot your device:

```
WAE# reload
```

Rebooting is optional; however, you might want to reboot to ensure that the boot flags are reset, and to ensure that subsequent console administrator logins do not bypass the password check.

Note In the WAAS software, the bootflags are reset to 0x0 on every reboot.

Recovering from Missing Disk-Based Software

This section describes how to recover from the following types of disk drive issues:

- Your WAAS device contains a single disk drive that needs to be replaced due to a disk failure.
- Your WAAS device contains two disk drives and you intentionally deleted the disk partitions on both drives (diks00 and disk01).

  Systems with two or more disk drives are normally protected automatically by RAID-1 on critical system partitions, so the procedures in this section do not need to be followed when replacing a disk drive in a multi-drive system.

To recover from this condition, follow these steps:

Step 1 Deactivate the device by completing the following steps:

a. From the WAAS Central Manager menu, go to **Devices >** *device-name*.

b. Choose *device-name* **> Activation**. The Device Activation window appears.

c. Uncheck the **Activate** check box, and then click **Submit**.

   The device is deactivated.

Step 2 Power down the device and replace the failed hard drive.

Step 3 Power on the device.

Step 4 Install the WAAS software. For more information on initial configuration, see the *Cisco Wide Area Application Services Quick Configuration Guide*.

Step 5 Use the CMS identity recovery procedure to recover the device CMS identity and associate this device with the existing device record on the WAAS Central Manager. For more information, see the "Recovering WAAS Device Registration Information" section on page 16-23.

Recovering WAAS Device Registration Information

Device registration information is stored both on the device itself and on the WAAS Central Manager. If a device loses its registration identity or needs to be replaced because of a hardware failure, the WAAS network administrator can issue a CLI command to recover the lost information, or in the case of adding a new device, assume the identity of the failed device.

To recover lost registration information, or to replace a failed device with a new one having the same registration information, follow these steps:

Step 1 Mark the failed device as "Inactive" and "Replaceable" by completing the following steps:

a. From the Central Manager menu, choose **Devices >** *device-name*.

b. Choose *device-name* **> Activation**.

c. Uncheck the **Activate** check box. The window refreshes, displaying a check box for marking the device as replaceable.

d. Check the **Replaceable** check box, and click **Submit**.

Note This check box appears in the GUI only when the device is inactive.

Step 2 If the failed device is configured as a non-optimizing peer with another device, disable the peer settings on the other device.

A pop-up message appears if the failed device is a non-optimizing peer and the message indicates the device that is its non-optimizing peer. When a device is replaced, its device ID changes and so the non-optimizing peer configuration must be updated.

a. From the WAAS Central Manager menu, choose **Configure > Global > Peer Settings**. The Peer Settings window for all devices appears.

b. Click the **Edit** icon next to the non-optimizing peer device identified in the message, which will appear in red because its peer is unknown. The Peer Settings window for that device appears.

c. Click the **Remove Device Settings** icon in the taskbar.

d. Click **Submit**.

Step 3 Configure a system device recovery key as follows:

a. From the WAAS Central Manager menu, choose **Configure > Global > System Properties**.

b. Click the **Edit** icon next to the System.device.recovery.key property. The Modifying Config Property window appears.

c. Enter a password in the Value field, and click **Submit**. The default password is **default**.

Step 4 Configure the basic network settings for the new device.

Step 5 Open a Telnet session to the device CLI and enter the **cms recover identity** *keyword* EXEC command. *keyword* is the device recovery key that you configured in the WAAS Central Manager GUI.

When the WAAS Central Manager receives the recovery request from the WAAS device, it searches its database for the device record that meets the following criteria:

- The record is inactive and replaceable.
- The record has the same hostname or primary IP address as given in the recovery request.

If the recovery request matches the device record, then the WAAS Central Manager updates the existing record and sends the requesting device a registration response. The replaceable state is cleared so that no other device can assume the same identity. When the WAAS device receives its recovered registration information, it writes it to file, initializes its database tables, and starts.

Step 6 Enter the following commands to enable the CMS service on the device:

```
WAE# config
WAE(config)# cms enable
WAE(config)# exit
```

Step 7 Activate the device:

a. From the WAAS Central Manager menu, choose **Devices >** *device-name*.

b. Choose *Device Name* **> Activation**. The WAAS device status should be Online.

c. Check the **Activate** check box, and click **Submit**.

Step 8 Reconfigure the device peer settings, if the device was configured as a non-optimizing peer with another device (see the "Information About Clustering Inline WAEs" section on page 5-53).

Step 9 Save the device configuration settings by entering the **copy running-config startup-config** EXEC command.

Performing Disk Maintenance for RAID-1 Systems

WAAS supports hot-swap functionality for both failed disk replacement and scheduled disk maintenance. When a disk fails, WAAS automatically detects the disk failure, marks the disk as bad, and removes the disk from the RAID-1 volume. To schedule disk maintenance, you must manually shut down the disk.

You must wait for the disk to be completely shut down before you physically remove the disk from the WAE. When the RAID removal process is complete, WAAS generates a disk failure alarm and trap. In addition, a syslog ERROR message is logged.

Note If the removal event (such as, a disk failure or software shutdown) occurs while the RAID array is in the rebuild process, the RAID removal process may take up to 1 minute to complete. The duration of this process depends on the size of the disk.

If the WAAS software removes a failed disk during the RAID rebuild process, a RAID rebuild failure alarm is generated. If you administratively shut down the disk during the RAID rebuild process, a RAID rebuild abort alarm is generated instead.

When you install a replacement disk, the WAAS software detects the replacement disk and performs compatibility checks on the disk, initializes the disk by creating partitions, and adds the disk to the software RAID to start the RAID rebuild process.

If the newly inserted disk has the same disk ID as a disk that was previously marked bad in the same physical slot, then the disk will not be mounted, and the post-replacement checks, initialization, and RAID rebuilding will not occur.

A newly installed disk must be of the same type and speed as the old disk and it must meet the following compatibility requirements:

- If the replacement disk is for disk00, disk02, or disk04 of a RAID pair, the replacement disk must be the same size as the running disk in the array.
- If the replacement disk is for disk01, disk03, or disk05 of a RAID pair, then the replacement disk must have the same or greater RAID capacity as the running disk in the array.

Compatibility checks, which are part of the hot-swap process, check for capacity compatibility. Incompatibility generates an alarm and aborts the hot-swap process.

To perform disk maintenance, follow these steps:

Step 1 Manually shut down the disk.

a. Enter global configuration mode and then enter the **disk disk-name** *diskxx* **shutdown** command:

```
WAE# configure
WAE(config)# disk disk-name diskxx shutdown
```

b. Wait for the disk to be completely shut down before you physically remove the disk from the WAE. When the RAID removal process is complete, WAAS generates a disk failure alarm and trap. In addition, a syslog ERROR message is logged.

Note We recommend that you disable the **disk error-handling reload** option, if enabled, because it is not necessary to power down the system to remove a disk.

Step 2 Insert a replacement disk into the slot in the WAE. The replacement disk must have a disk ID number that is different from the disk that it is replacing.

Step 3 Reenable the disk by entering the **no disk disk-name** *diskxx* **shutdown** global configuration command.

Replacing Disks in RAID-5 Systems

To remove and replace a physical disk drive in a system that uses a RAID-5 logical drive, follow these steps:

Step 1 Enter the **disk disk-name** *diskxx* **replace** command in EXEC mode at the WAAS CLI on the WAE.

Step 2 Verify that the disk drive *diskxx* is in the Defunct state by entering the **show disks details** command in EXEC mode. The RAID logical drive is in the Critical state at this point.

Step 3 Move the handle on the drive to the open position (perpendicular to the drive).

Step 4 Pull the hot-swap drive assembly from the bay.

Step 5 Wait for one minute and then insert the new drive into the same slot by aligning the replacement drive assembly with guide rails in the bay and sliding the drive assembly into the bay until it stops. Make sure that the drive is properly seated in the bay.

Step 6 Close the drive handle.

Step 7 Check the hard disk drive status LED to verify that the hard disk drive is operating correctly. If the amber hard disk drive status LED for a drive is lit continuously, that drive is faulty and must be replaced. If the green hard disk drive activity LED is flashing, the drive is being accessed.

Step 8 Wait 1 minute and then verify that the replaced disk drive is in the Rebuilding state by using the **show disks details** command in EXEC mode.

Note The ServeRAID controller automatically starts the rebuild operation when it detects the removal and reinsertion of a drive that is part of the logical RAID drive.

Step 9 Wait until the rebuild operation is complete. You can check if the rebuild operation is complete by using the **show disks details** command in EXEC mode. The physical drive state will be Online and the RAID logical drive state will be Okay after the rebuild operation is completed.

A 300-GB SAS drive may take up to 5 hours to finish rebuilding.

If you have multiple disk failures and your RAID-5 logical status is Offline, you must recreate the RAID-5 array by following these steps:

Step 1 Enter global configuration mode and then enter the **disk logical shutdown** command to disable the RAID-5 array.

Step 2 Enter the **write** command in EXEC mode to save the running configuration to NVRAM.

Step 3 Enter the **reload** command in EXEC mode to reload the system.

Step 4 Enter the **show disks details** command in EXEC mode to check the system configuration after the system is rebooted. At this point, the disks are not mounted and the logical RAID drive should be in the Shutdown state.

Step 5 Enter the **disk recreate-raid** command in EXEC mode to recreate the RAID-5 array.

Step 6 After successful execution of the previous command, enter global configuration mode and then enter the **no disk logical shutdown** command to disable the logical disk shutdown configuration.

Step 7 Enter the **write** command in EXEC mode to save the configuration to NVRAM.

Step 8 Enter the **reload** command in EXEC mode to reload the system.

Step 9 Enter the **show disks details** command in EXEC mode to check the system configuration after the system is rebooted. At this point, the disks should be mounted and the logical RAID drive should not be in the Shutdown state.

Step 10 Wait until the rebuild operation is complete. You can check if the rebuild operation is complete by using the **show disks details** command in EXEC mode. The physical drive state will be Online and the RAID logical drive state will be Okay after the rebuild operation is completed.

It takes several hours to finish rebuilding the RAID-5 array.

After a multiple disk failure or RAID controller failure, and after the drives are replaced and the RAID disk is rebuilt, the logical disk may remain in the error state. To reenable the disk, use the **no disk logical shutdown force** command, then reload the WAE.

Configuring the Central Manager Role

The WAAS software implements a standby WAAS Central Manager. This process allows you to maintain a copy of the WAAS network configuration on a second WAAS Central Manager device. If the primary WAAS Central Manager fails, the standby can be used to replace the primary.

For interoperability, when a standby WAAS Central Manager is used, it must be at the same software version as the primary WAAS Central Manager to maintain the full WAAS Central Manager configuration. Otherwise, the standby WAAS Central Manager detects this status and does not process any configuration updates that it receives from the primary WAAS Central Manager until the problem is corrected.

Note Primary and standby Central Managers communicate on port 8443. If your network includes a firewall between primary and standby Central Managers, you must configure the firewall to allow traffic on port 8443 so that the Central Managers can communicate and stay synchronized.

This section contains the following topics:

Converting a WAE to a Standby Central Manager

This section describes how to convert a WAE that is operating as an application accelerator to a standby Central Manager.

There are two types of WAAS software files, as follows:

- Universal—Includes Central Manager, Application Accelerator, and AppNav Controller functionality.
- Accelerator only—Includes Application Accelerator and AppNav Controller functionality only. If you want to change an Application Accelerator or AppNav Controller to a Central Manager, you must use the Universal software file.

If the WAE is operating with an Accelerator only image, you cannot convert it to a Central Manager until after you update it with the Universal software file, reload the device, change the device mode to central-manager, and then reload the device again. For information on updating a WAE, see the "Upgrading the WAAS Software" section on page 16-1.

You can use the **show version** EXEC command to check if the WAE is running an Accelerator only image. Also, the **show running-config** EXEC command displays the image type.

To convert a WAE with a Universal image to a standby Central Manager, follow these steps:

Step 1 Deregister the WAE from the Central Manager using the **cms** command:

```
WAE# cms deregister force
```

This command cleans up any previous association to any other Central Manager.

Step 2 Configure the device mode as Central Manager using the **device** command:

```
WAE# configure
WAE(config)# device mode central-manager
```

You must reload the device to apply the changes.

Step 3 Configure the Central Manager role as standby using the **central-manager** command:

```
WAE(config)# central-manager role standby
```

Step 4 Configure the address of primary Central Manager using the **central-manager** command:

```
WAE(config)# central-manager address cm-primary-address
```

Step 5 Enable the CMS service using the **cms** command:

```
WAE(config)# cms enable
```

Converting a Primary Central Manager to a Standby Central Manager

To convert a primary Central Manager to a standby Central Manager, follow these steps:

Step 1 Deregister the Central Manager using the **cms** command:

```
WAE# cms deregister
```

This command cleans up any previous association to any other Central Manager.

Step 2 Configure the Central Manager role as standby using the **central-manager** command:

```
WAE# configure
WAE(config)# central-manager role standby
```

Step 3 Configure the address of primary Central Manager using the **central-manager** command:

```
WAE(config)# central-manager address cm-primary-address
```

Step 4 Enable the CMS service using the **cms** command:

```
WAE(config)# cms enable
```

Converting a Standby Central Manager to a Primary Central Manager

If your primary WAAS Central Manager becomes inoperable, you can manually reconfigure one of your warm standby Central Managers to be the primary Central Manager. Configure the new role by using the global configuration **central-manager role primary** command as follows:

```
WAE# configure
WAE(config)# central-manager role primary
```

This command changes the role from standby to primary and restarts the management service to recognize the change.

If the previous failed primary Central Manager again becomes available, you can recover it to again become the primary Central Manager. For details, see the "Central Manager Failover and Recovery" section on page 16-30.

If you switch a warm standby Central Manager to primary while your primary Central Manager is still online and active, both Central Managers detect each other, automatically shut themselves down, and disable management services. The Central Managers are switched to halted, which is automatically saved in flash memory.

To return halted WAAS Central Managers to an online status, decide which Central Manager should be the primary device and which should be the standby device. On the primary device, execute the following CLI commands:

```
WAE# configure
WAE(config)# central-manager role primary
WAE(config)# cms enable
```

On the standby device, execute the following CLI commands:

```
WAE# configure
WAE(config)# central-manager role standby
WAE(config)# central-manager address cm-primary-address
WAE(config)# cms enable
```

Switching Both Central Manager Roles

Caution When you switch a WAAS Central Manager from primary to standby, the configuration on the Central Manager is erased. The Central Manager, after becoming a standby, will begin replicating its configuration information from whichever Central Manager is now the primary. If standby and primary units are not synchronized before switching roles, important configuration information can be lost.

Before you switch Central Manager roles, follow these steps:

Step 1 Ensure that your Central Manager devices are running the same version of WAAS software.

Step 2 Synchronize the physical clocks on both devices so that both WAAS Central Managers have the same Coordinated Universal Time (UTC) configured.

Step 3 Ensure that the standby is synchronized with the primary by checking the status of the following items:

a. Check the online status of your devices.

The original standby Central Manager and all currently active devices should be showing as online in the Central Manager GUI. This step ensures that all other devices know about both Central Managers.

b. Check the status of recent updates from the primary WAAS Central Manager.

Use the **show cms info** EXEC command and check the time of the last update. To be current, the value of the Time of last config-sync field should be between 1 and 5 minutes old. You are verifying that the standby WAAS Central Manager has fully replicated the primary WAAS Central Manager configuration.

If the update time is not current, determine whether or not there is a connectivity problem or if the primary WAAS Central Manager is down. Fix the problem, if necessary, and wait until the configuration has replicated, as indicated by the time of the last update.

Step 4 Switch roles in the following order:

a. Switch the original primary to standby mode:

```
WAE1# configure
WAE1(config)# central-manager role standby
WAE(config)# cms enable
```

b. Switch the original standby to primary mode:

```
WAE2# configure
WAE2(config)# central-manager role primary
WAE(config)# cms enable
```

The CMS service is restarted automatically when you configure a role change.

Central Manager Failover and Recovery

If your primary WAAS Central Manager becomes inoperable, you can reconfigure one of your standby Central Managers to be the primary Central Manager and then later, when the failed Central Manager becomes available, you can reconfigure it to be primary again. Follow these steps:

Step 1 Convert a standby Central Manager to be the primary Central Manager as described in the "Converting a Standby Central Manager to a Primary Central Manager" section on page 16-28.

Step 2 When the failed Central Manager is again available, configure it as a standby Central Manager as described in the "Converting a Primary Central Manager to a Standby Central Manager" section on page 16-28, beginning with Step 2. Skip the first step and do not use the **cms deregister** command.

Step 3 Switch both Central Manager roles as described in the "Switching Both Central Manager Roles" section on page 16-29.

Enabling Disk Encryption

Disk encryption addresses the need to securely protect sensitive information that flows through deployed WAAS systems and that is stored in WAAS persistent storage. The disk encryption feature includes two aspects: the actual data encryption on the WAE disk and the encryption key storage and management.

When you enable disk encryption, all data in WAAS persistent storage will be encrypted. The encryption key for unlocking the encrypted data is stored on the Central Manager, and key management is handled by the Central Manager. When you reboot the WAE after configuring disk encryption, the WAE retrieves the key from the Central Manager automatically, allowing normal access to the data that is stored in WAAS persistent storage.

Note If a WAE is unable to reach the WAAS Central Manager during a reboot, it will do everything except mount the encrypted partitions. In this state, all traffic will be handled as pass-through. Once communication with the WAAS Central Manager is restored (and the encryption key is obtained), the encrypted partitions are mounted. There is no loss of cache content.

Disk encryption requirements are as follows:

- You must have a Central Manager configured for use in your network.
- Your WAE devices must be registered with the Central Manager.
- Your WAE devices must be online (have an active connection) with the Central Manager. This requirement applies only if you are enabling disk encryption.
- You must reboot your WAE for the disk encryption configuration to take effect.

After you reboot your WAE, the encryption partitions are created using the new key, and any previously existing data is removed from the partition.

Any change to the disk encryption configuration, whether to enable or disable encryption, causes the disk to clear its cache. This feature protects sensitive customer data from being decrypted and accessed should the WAE ever be stolen.

If you enable disk encryption and then downgrade to a software version that does not support this feature, you will not be able to use the data partitions. In such cases, you must delete the disk partitions after you downgrade.

To enable and disable disk encryption from the Central Manager GUI, choose **Devices >** *device-name*, then choose **Configure > Storage > Disk Encryption**. To enable disk encryption, check the **Enable** check box and click **Submit**. This box is unchecked by default. To disable disk encryption, uncheck the **Enable** check box and click **Submit**.

To enable and disable disk encryption from the WAE CLI, use the **disk encrypt** global configuration command.

Note If you are using an NPE image, the disk encryption feature has been disabled for use in countries where disk encryption is not permitted.

When you enable or disable disk encryption, the file system is reinitialized during the first subsequent reboot. Reinitialization may take from ten minutes to several hours, depending on the size of the disk partitions. During this time, the WAE will be accessible, but it will not be providing any services.

If you change the Central Manager IP address, or if you relocate the Central Manager, or replace one Central Manager with another Central Manager that has not copied over all of the information from the original Central Manager, and you reload the WAE when disk encryption is enabled, the WAE file system will not be able to complete the reinitialization process or obtain the encryption key from the Central Manager.

If the WAE fails to obtain the encryption key, disable disk encryption by using the **no disk encrypt enable** global configuration command from the CLI, and reload the WAE. Ensure connectivity to the Central Manager before you enable disk encryption and reload the WAE. This process will clear the disk cache.

Note When a standby Central Manager has been in service for at least 2 times the datafeed poll rate time interval (approximately 10 minutes) and has received management updates from the primary Central Manager, the updates will include the latest version of the encryption key. Failover to the standby in this situation occurs transparently to the WAE. The datefeed poll rate defines the interval for the WAE to poll the Central Manager for configuration changes. This interval is 300 seconds by default.

To view the encryption status details, use the **show disks details** EXEC command. While the file system is initializing, **show disks details** displays the following message: "`System initialization is not finished, please wait...`" You may also view the disk encryption status, whether it is enabled or disabled, in the Central Manager GUI, Device Dashboard window.

Configuring a Disk Error-Handling Method

Note Configuring and enabling disk error handling is no longer necessary for devices that support disk hot-swap. In WAAS 4.0.13 and later, the software automatically removes from service any disk with a critical error.

The WAAS Central Manager allows you to configure how disk errors should be handled and to define a disk device error-handling threshold for WAAS devices running software versions 4.1.1 and earlier.

If the bad disk drive is a critical disk drive, and the automatic reload feature is enabled, then the WAAS software marks the disk drive "bad" and the WAAS device is automatically reloaded. After the WAAS device is reloaded, a syslog message and an SNMP trap are generated.

Note The automatic reload feature is automatically enabled and not configurable on devices running WAAS version 4.1.3 and later.

To configure a disk error-handling method using the WAAS Central Manager GUI, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Storage > Disk Error Handling**.

The Disk Error Handling Settings window appears.

Step 3 Check the **Enable** check box to enable the window for configuration, and then check the following options as necessary:

- **Enable Disk Error Handling Reload**—Forces the device to reload the disk if the file system (sysfs) (disk00) has problems. This option appears only for device groups and version 4.1.1 and earlier WAAS devices and is disabled by default. It does not apply to WAAS versions later than 4.1.1.
- **Enable Disk Error Handling Remap**—Forces the disks to attempt to remap disk errors automatically. This option is enabled by default.
- **Enable Disk Error Handling Threshold**—Specifies the number of disk errors allowed before the disk is marked as bad. You must enter a number between 0 to 100 in the Threshold field. The default threshold is 10. This option is disabled by default. This option appears only for device groups and version 4.1.1 and earlier WAAS devices. It does not apply to WAAS versions later than 4.1.1.

Step 4 Click **Submit** to save the settings.

Enabling Extended Object Cache

The WAAS Central Manager allows you to configure additional disk space for CIFS object caching. Disk caching is used for DRE and CIFS to serve content at the edge. The extended object cache feature lets you choose the amount of disk storage that is used by CIFS and Virtual Blade Services.

This feature is supported only on WAVE-694-16G, WAVE-694-24G, WAE-674-4G, or WAE-674-8G models and does not have any effect on other models in a device group.

Note If extended object cache is enabled and the device is downgraded to a version prior to 4.2.1, all CIFS cache data, DRE cache data, and virtual blade data is lost.

To enable extended object cache using the WAAS Central Manager GUI, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Storage > Extended Object Cache**.

The Extended Disk Space Settings window appears.

Step 3 Check the **Enable** check box to enable extended object cache. See Table 16-3 and Table 16-4 for sizing details.

Note If extended object cache is enabled on WAVE-694-16G, WAVE-694-24G, WAE-674-4G, or WAE-674-8G models with a virtual blade disk size that is greater than 30 GB, the device goes into override mode. In override mode, a Force Device Group settings button will be present on the device page (or device group page).

If you have a virtual blade enabled with a disk size that is greater than 30 GB, you must stop the virtual blade and remove the configuration before enabling extended object cache. Otherwise, the size of the virtual blade filesystem will be reduced to 30 GB.

Step 4 Click **Submit** to save the settings.

The disk cache size varies depending upon the features enabled.

Table 16-3 shows the usable disk cache sizes for the WAE-674-4G platform.

Table 16-3 WAE-674-4G Platform Disk Cache Sizes

	Extended Object Cache Disabled		Extended Object Cache Enabled	
Disk Partition	**Virtual Blade Disabled**	**Virtual Blade Enabled**	**Virtual Blade Disabled**	**Virtual Blade Enabled**
DRE Cache	114 GB	120 GB	114 GB	114 GB
CIFS Object Cache	96 GB	96 GB	340 GB	305 GB
Virtual Blade	--	242 GB	--	29 GB

Table 16-4 shows the usable disk cache sizes for the WAE-674-8G platform.

Table 16-4 WAE-674-8G Platform Disk Cache Sizes

Disk Partition	Extended Object Cache Disabled		Extended Object Cache Enabled	
	Virtual Blade Disabled	Virtual Blade Enabled	Virtual Blade Disabled	Virtual Blade Enabled
DRE Cache	304 GB	143 GB	143 GB	143 GB
CIFS Object Cache	96 GB	96 GB	310 GB	275 GB
Virtual Blade	--	210 GB	--	27 GB

Table 16-5 shows the usable disk cache sizes for the WAVE-694-16G platform.

Table 16-5 WAVE-694-16G Platform Disk Cache Sizes

Disk Partition	Extended Object Cache Disabled and Virtual Blade Enabled[1]	Extended Object Cache and Virtual Blade Enabled
DRE Cache	114 GB	114 GB
CIFS Object Cache	77 GB	304 GB
Virtual Blade	161 GB	22 GB

1. The Virtual Blade cannot be disabled on this platform.

Table 16-6 shows the usable disk cache sizes for the WAVE-694-24G platform.

Table 16-6 WAVE-694-24G Platform Disk Cache Sizes

Disk Partition	Extended Object Cache Disabled and Virtual Blade Enabled[1]	Extended Object Cache and Virtual Blade Enabled
DRE Cache	190 GB	142 GB
CIFS Object Cache	77 GB	272 GB
Virtual Blade	161 GB	22 GB

1. The Virtual Blade cannot be disabled on this platform.

Activating All Inactive WAAS Devices

To activate all inactivated WAAS devices in your network, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices > All Devices**. The All Devices window appears.

Step 2 Click the **Activate all inactive WAEs** icon in the taskbar. The Activate All Inactive WAEs window appears.

Step 3 Choose an existing location for all inactivated WAAS devices by clicking the **Select an existing location for all inactive WAEs** radio button, and then choose a location from the drop-down list.

Alternatively, choose to create a new location for each inactive device by clicking the **Create a new location for each inactive WAE** radio button. Specify a parent location for all newly created locations by choosing a location from the Select a parent location for all newly created locations drop-down list.

Step 4 Click **Submit**. The inactive WAEs are reactivated and placed in the specified location.

Rebooting a Device or Device Group

Using the WAAS Central Manager GUI, you can reboot a device or device group remotely.

To reboot an individual device, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*. The device Dashboard appears.

Step 2 Click the **Reload** icon in the Device Info pane. You are prompted to confirm your decision.

Step 3 Click **OK** to confirm that you want to reboot the device.

To reboot a device from the CLI, use the **reload** EXEC command.

If you reboot a WAAS Central Manager that has the secure store enabled with user-provided passphrase mode, you must reopen the secure store after the reboot by using the **cms secure-store open** EXEC command.

To reboot an entire device group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Device Groups >** *device-group-name*. The Modifying Device Group window appears.

Step 2 In the taskbar, click the **Reboot All Devices in Device Group** icon. You are prompted to confirm your decision.

Step 3 Click **OK** to confirm that you want to reboot the device group.

Performing a Controlled Shutdown

A controlled shutdown refers to the process of properly shutting down a WAAS device without turning off the power on the device (the fans continue to run and the power LED remains on). With a controlled shutdown, all of the application activities and the operating system are properly stopped on the appliance, but the power remains on. Controlled shutdowns can help you minimize the downtime when the appliance is being serviced.

Caution If a controlled shutdown is not performed, the WAAS file system can be corrupted. It also takes longer to reboot the appliance if it was not properly shut down.

You can perform a controlled shutdown from the CLI by using the **shutdown** EXEC command. For more details, see the *Cisco Wide Area Application Services Command Reference*.

If you are running WAAS on a network module that is installed in a Cisco access router, perform a controlled shutdown from the router CLI by using the **service-module integrated-service-engine** *slot/unit* **shutdown** EXEC command. For more details, see the document *Configuring Cisco WAAS Network Modules for Cisco Access Routers*.

CHAPTER 17

Monitoring and Troubleshooting Your WAAS Network

This chapter describes the monitoring and troubleshooting tools available in the WAAS Central Manager GUI that can help you identify and resolve issues with your WAAS system.

For additional advanced WAAS troubleshooting information, see the *Cisco WAAS Troubleshooting Guide for Release 4.1.3 and Later* on Cisco DocWiki.

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

This chapter contains the following sections:

Viewing System Information from the System Dashboard Window

The WAAS Central Manager GUI allows you to view general and detailed information about your WAAS network from the System Dashboard window. This section describes the System Dashboard window and contains the following topics:

- Alarm Panel, page 17-3
- Device Alarms, page 17-4

Figure 17-1 shows the System Dashboard window.

Figure 17-1 System Dashboard Window

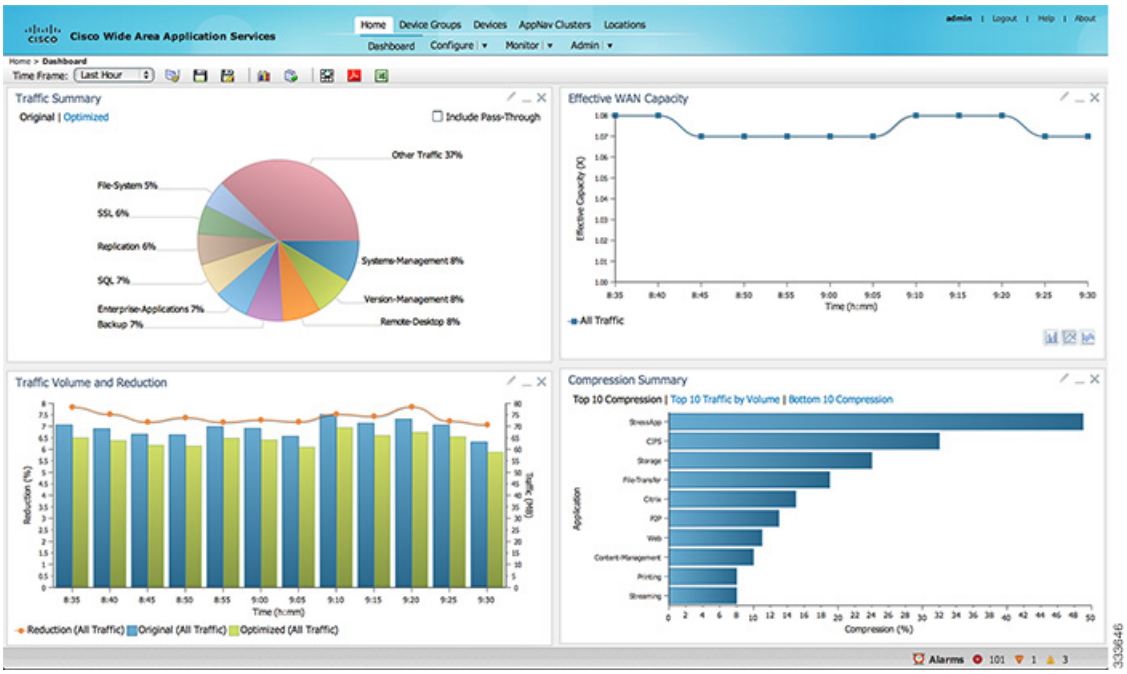

The information displayed in the charts in the System Dashboard window is based on a snapshot of your WAAS network that represents the state of your WAE devices at the end of every two polling periods. You may configure the interval between polls in the WAAS Central Manager GUI (**Configure > Global > System Properties > System.monitoring.collectRate**). The default polling rate is 300 seconds (5 minutes). Alarms are presented in real time and are independent of the polling rate.

Monitoring Graphs and Charts

The default System Dashboard window contains the following graphical displays about the application traffic processed by your WAAS system:

- Traffic Summary chart—Displays the applications with the highest percentage of traffic in the WAAS network for the last hour.
- Effective WAN Capacity graph—Displays the effective increased bandwidth capacity of the WAN link as a result of WAAS optimization, as a multiple of the actual bandwidth.
- Traffic Volume and Reduction graph—Displays the original and optimized traffic volume and percentage of traffic reduction over the last hour.

- Compression Summary chart—Displays the ten applications with the highest percentage of traffic reduction for the WAAS network for the last hour. The percent calculation excludes pass-through traffic.

Numbers shown in charts and graphs are rounded to whole units (KB, MB, or GB), while those displayed in tables are rounded to three decimal places. Data values exported to CSV files are in bytes, so are not rounded.

You can customize the graphical displays and tables that are displayed on the system dashboard. For more information, see the "Customizing a Dashboard or Report" section on page 17-10. Individual charts are described in more detail in the "Chart and Table Descriptions" section on page 17-14.

Much of the device, statistical, and alarm information that is presented in the system dashboard and associated graphs and charts is also available programmatically through the monitoring API. For more information, see the *Cisco Wide Area Application Services API Reference*.

Note You must synchronize the clock on each WAE device within 5 minutes of the primary and secondary WAAS Central Managers for statistics to be consistent and reliable. For information on using an NTP server to keep all your WAAS devices synchronized, see the "Configuring NTP Settings" section on page 10-5. Additionally, if the network delay for the Central Manager to receive statistical updates from the WAEs is greater than 5 minutes, statistics aggregation may not operate as expected.

Alarm Panel

The alarm panel provides a near real-time view of incoming alarms and refreshes every two minutes to reflect updates to the system alarm database.

To view the alarms panel, click **Alarms** at the bottom right side of the Central Manager window.

Only Active alarms can be acknowledged in the alarm panel. Pending, Offline, and Inactive alarms cannot be acknowledged in the alarm panel.

The alarm panel also allows you to filter your view of the alarms in the list. Filtering allows you to find alarms in the list that match the criteria that you set.

Figure 17-2 shows the alarm panel.

Figure 17-2 Alarm Panel

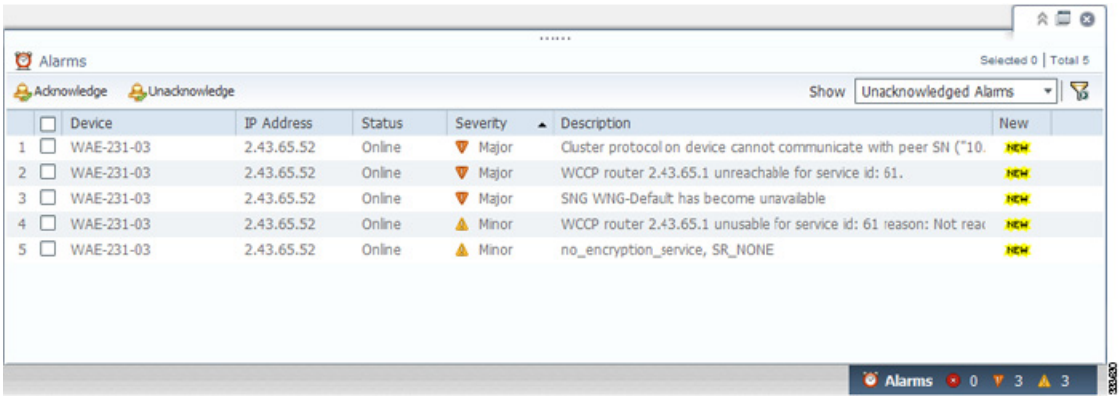

To acknowledge an active alarm, follow these steps:

Step 1 In the alarm panel, check the check box next to the name of the alarm that you want to acknowledge.

Step 2 Click the **Acknowledge** taskbar icon.

A dialog box pops up that allows you to enter comments about the alarm.

Step 3 Enter a comment and click **OK**. Alternatively, click **Cancel** to return to the alarm panel without completing the acknowledge action.

Comments enable you to share information about the cause or solution of a particular problem that caused the alarm. The comments field accepts up to 512 characters. You may use any combination of alpha, numeric, and special characters in this field.

To filter and sort alarms displayed in the alarm panel, follow these steps:

Step 1 From the Show drop-down list, choose one of the following filtering options:

- **All**
- **Quick Filter**
- **Unacknowledged Alarms**
- **Acknowledged Alarms**
- **Alarms for** *device-name* (shown in the device context)

Step 2 If you chose Quick Filter, enter match criteria in one or more fields above the list.

Step 3 To sort alarm entries, click a column header.

Entries are sorted alphabetically (in ASCII order). The sort order (ascending or descending) is indicated by an arrow in the column header that points up for ascending order.

Step 4 Choose **All** to clear the filter.

Device Alarms

Device alarms are associated with device objects and pertain to applications and services running on your WAAS devices. Device alarms are defined by the reporting application or service. Device alarms can also reflect reporting problems between the device and the WAAS Central Manager GUI. Table 17-1 describes the various device alarms that can appear.

Table 17-1 Device Alarms for Reporting Problems

Alarm	Alarm Severity	Device Status	Description
Device is offline	Critical	Offline	The device has failed to communicate with the WAAS Central Manager.
Device is pending	Major	Pending	The device status cannot be determined. This status can appear after a new device is registered but before the first configuration synchronization has been done.
Device is inactive	Minor	Inactive	The device has not yet been activated or accepted by the WAAS Central Manager.
Device has lower software version	Minor	Online	The device has an earlier software version than the WAAS Central Manager and it may not support some features.

Troubleshooting Devices Using Alerts

The WAAS Central Manager GUI allows you to view the alarms on each device and troubleshoot a device in the Troubleshooting Devices window.

To troubleshoot a device from the Troubleshooting Devices window, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices > All Devices** and click the device alarm light bar in the Device Status column to view alarms on a single device.

The Troubleshooting Devices window appears, either in the WAAS Central Manager window or as a separate popup window. (See Figure 17-3.)

Figure 17-3 Troubleshooting Devices Window

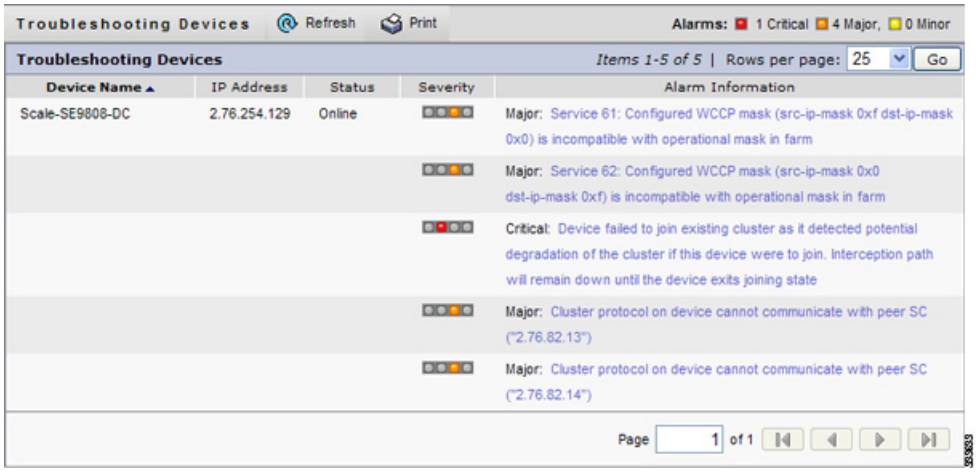

Step 2 In the Alarm Information column, hold your mouse over an alarm message until the Troubleshooting tools contextual menu appears. The popup menu provides links to the troubleshooting and monitoring windows in the WAAS Central Manager GUI.

Step 3 Choose the troubleshooting tool that you want to use, and click the link. The link takes you to the appropriate window in the WAAS Central Manager GUI. Table 17-2 describes the tools available for device alarms.

You can view the Troubleshooting Devices window for all devices by choosing **Monitor > Troubleshoot > Alerts** from the global context.

Table 17-2 Troubleshooting Tools for Device Alarms

Item	Navigation	Description
Update Software	Choose device, **Admin > Versioning > Software Update**	Displays Software Update window for this device. Appears only if the device software version is lower than the Central Manager.
Edit/Monitor Device	Device Dashboard	Displays Device Dashboard window for configuration.
Telnet to Device	Opens a Telnet window	Initiates a Telnet session using the device IP address.
View Device Log	Choose device, **Admin > History > Logs**	Displays system message logs filtered for this device.
Run Show Commands	Choose device, **Monitor > CLI Commands > Show Commands**	Displays the device **show** command tool. For more information, see the "Using the show and clear Commands from the WAAS Central Manager GUI" section on page 17-61.

Viewing Device Information

The WAAS Central Manager GUI allows you to view basic and detailed information about a device from the following two windows:

- Devices Window—Displays a list of all the devices in your WAAS network along with basic information about each device such as the device status and the current software version installed on the device.
- Device Dashboard Window—Displays detailed information about a specific device, such as the installed software version and whether the device is online or offline.

Each window is explained in the sections that follow.

Devices Window

The Devices window lists all the WAAS devices that are registered with the WAAS Central Manager. To view this list, choose **Devices > All Devices** in the WAAS Central Manager GUI.

Figure 17-4 shows an example of the Devices window.

Figure 17-4 Devices Window

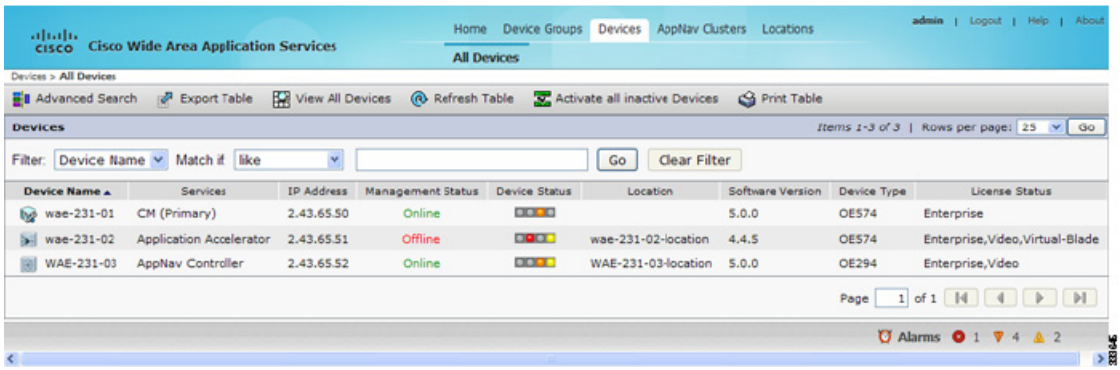

This window displays the following information about each device:

- Services enabled on the device. See Table 17-3 for a description of these services.
- IP address of the device.
- Management Status (Online, Offline, Pending, or Inactive). For more information about the status, see the "Device Alarms" section on page 17-4.
- Device Status. The system status reporting mechanism uses four alarm lights to identify problems that need to be resolved. Each light represents a different alarm level as follows:
  - Green—No alarms (the system is in excellent health)
  - Yellow—Minor alarms
  - Orange—Major alarms
  - Red—Critical alarms

  When you roll your mouse over the alarm light bar, a popup message provides further details about the number of alarms. Click the alarm light bar to troubleshoot the device. For more information, see the "Troubleshooting Devices Using Alerts" section on page 17-5.
- Location associated with the device. For more information about locations, see Chapter 3, "Using Device Groups and Device Locations." You can view reports that aggregate data from all devices in a location (see the "Location Level Reports" section on page 17-36).
- Software Version installed and running on the device. For WAAS Express devices, both the Cisco IOS and WAAS Express versions are shown.
- Device Type. If you see a type such as OE294, the numbers refer to the model number, such as the WAVE-294 in this example. NME-WAE refers to a NME-WAE module and SM-WAE refers to a SM-SRE module. For WAAS Express devices, the router platform is displayed. For vWAAS devices, OE-VWAAS is displayed.
- License Status. Displays the installed licenses. See Table 17-4 for a description of the possible values.

Any WAE devices that are at a higher software version level than the WAAS Central Manager are shown in red. Also, if the standby WAAS Central Manager has a different version level from the primary WAAS Central Manager, the standby WAAS Central Manager is shown in red.

You can filter your view of the devices in the list by using the Filter and Match If fields above the list. Enter a filter string in the text field and click the **Go** button to apply the filter. The filter settings are shown below the list. Click the **Clear Filter** button to clear the filter and show all devices. Filtering allows you to find devices in the list that match the criteria that you set.

Table 17-3 Service Descriptions

Service	Description
CM (Primary)	The device has been enabled as the primary WAAS Central Manager. For information on primary and standby Central Manager devices, see the "Converting a Standby Central Manager to a Primary Central Manager" section on page 16-28.
CM (Standby)	The device has been enabled as a standby WAAS Central Manager. For information on primary and standby Central Manager devices, see the "Converting a Standby Central Manager to a Primary Central Manager" section on page 16-28.
Application Accelerator	The device has been enabled as an application accelerator.
AppNav Controller	The device has been enabled as an AppNav Controller.
Replication Accelerator	The device has been enabled as a replication accelerator. (Supported only on 4.0.19 or later 4.0.x devices.)
WAAS Express	The device is an IOS router with the WAAS Express functionality enabled.

Table 17-4 License Status Descriptions

License Status	Description
Not Active	No license is installed or the first configuration synchronization has not yet happened.
Transport, Enterprise, Video, VB	The listed licenses are installed.
Active	A WAAS Express device is registered but the first configuration synchronization has not yet happened.
Permanent	A WAAS Express device has a permanent license installed.
Evaluation, Expires in X weeks Y days	A WAAS Express device has an evaluation license installed and it expires after the indicated period.
Expired	A WAAS Express device has an expired evaluation license. A permanent license must be obtained for this device to operate.
N/A	The license status is not applicable because the device version is 4.0.

Device Dashboard Window

The Device Dashboard window provides detailed information about a WAAS device such as the device model, IP address, interception method, and device-specific charts. (See Figure 17-5.)

To access the Device Dashboard window, choose **Devices >** *device-name*.

Figure 17-5 Device Dashboard Window

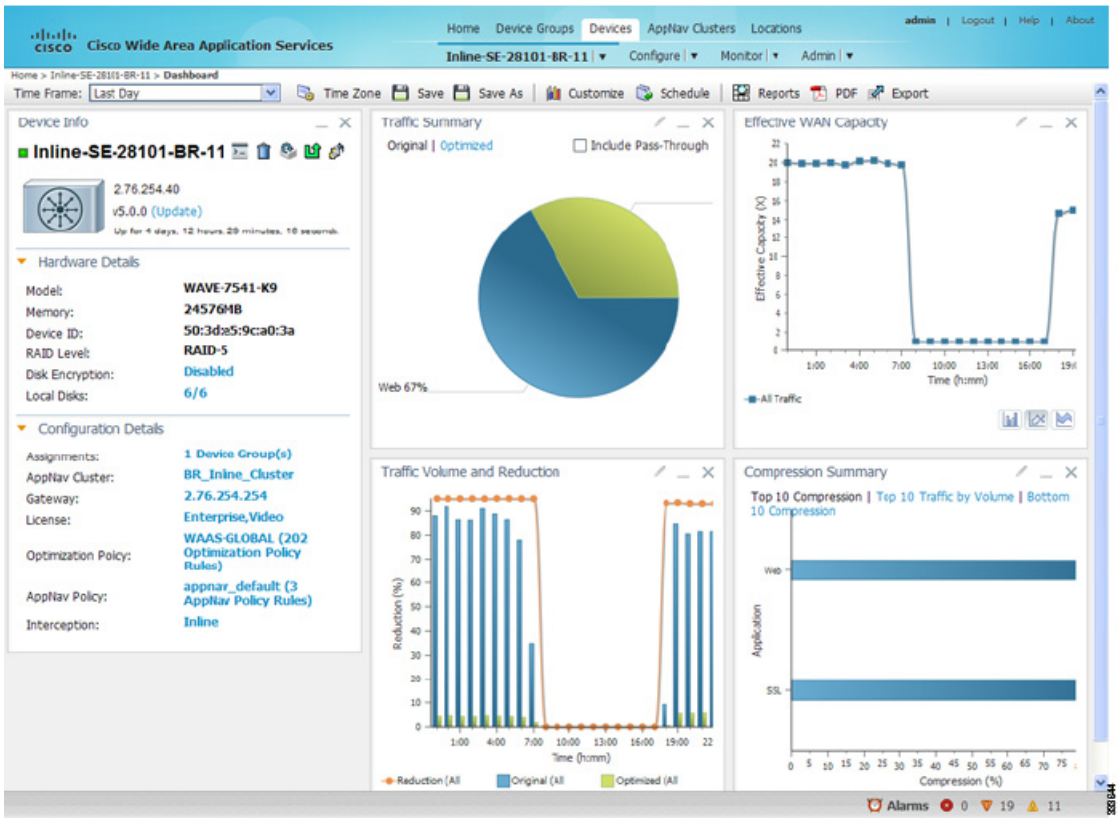

The Device Dashboard window for a WAAS Express device looks slightly different. It lacks some WAE-specific information and controls.

From the Device Dashboard window, you can perform the following tasks:

- View charts and graphs about the application traffic processed by the selected WAE device. (No charts or graphs are displayed if a WAAS Central Manager device is selected.)
- Customize the charts displayed in the window. For more information, see the "Customizing a Dashboard or Report" section on page 17-10. Individual charts are described in more detail in the "Chart and Table Descriptions" section on page 17-14.
- View basic details such as whether the device is online, the device's IP address and hostname, the software version running on the device, and the amount of memory installed in the device, the license status, and so forth.
- View the device groups to which the device belongs. For more information about device groups, see Chapter 3, "Using Device Groups and Device Locations."
- View the users that are defined on the device and unlock any locked out users. For more information, see the "Viewing and Unlocking Device Users" section on page 17-10. (Not available on WAAS Express devices.)
- Click the **Update** link to update the software on the device. For more information, see Chapter 16, "Maintaining Your WAAS System." (Not available on WAAS Express devices.)
- Click the **Telnet** icon to establish a Telnet session into the device and issue CLI commands.
- Click the **Delete Device** icon to delete the device.

- Click the **Full Update** icon to reapply the device configuration from the Central Manager to the device. (Not available on WAAS Express devices.)
- Click the **Reload** icon to reboot the device. (Not available on WAAS Express devices.)
- Click the **Restore Default Policies** icon to restore the default predefined policies on the device. For more information, see the "Restoring Optimization Policies and Class Maps" section on page 13-58.
- Assign and unassign the device to device groups. For more information, see Chapter 3, "Using Device Groups and Device Locations."
- For a WAAS Express device, a WAAS Enabled Interfaces item shows the number of interfaces on which WAAS optimization is enabled. You can click the number to go to the Network Interfaces configuration screen, which displays device interface details and allows you to enable or disable optimization on the available interfaces. For more details, see the "Configuring Optimization on WAAS Express Interfaces" section on page 6-15.
- For a WAAS Express device, you can view the DRE item to determine if the device supports data redundancy elimination (DRE) optimization, which is not supported on some WAAS Express device models. This item reads Supported or Unsupported.
- For a WAAS Express device, you can view the SSL item to determine if SSL acceleration is available. This item reads Available or Unavailable.
- For a vWAAS device, the No of CPUs, Max TCP Connections, and Interception Method fields are shown. If VPATH is enabled for the vWAAS device, it is indicated in the Interception Method field. For more details, see the "Configuring VPATH Interception on a vWAAS Device" section on page 5-55.

Viewing and Unlocking Device Users

To view the users that are defined on a WAAS device, go to **Devices >** *device-name*, and then from the *device-name* menu, choose **Device Users** (on a Central Manager device, choose **CM Users**).

The list of users is displayed in a table that shows the username, number of login failures, maximum number of login failures allowed, and the time of the last failed login. To view the details on a user, click the **View** icon next to the user.

If a user is locked out because they reached the maximum number of failed login attempts, you can unlock the user by checking the box next to the username and clicking the **Unlock** button below the table.

Customizing a Dashboard or Report

You can customize the system and device dashboards and any report in the same way. For more information about creating custom reports, see the "Managing Reports" section on page 17-43.

An example of a report is shown in Figure 17-6.

Figure 17-6 Report Pane

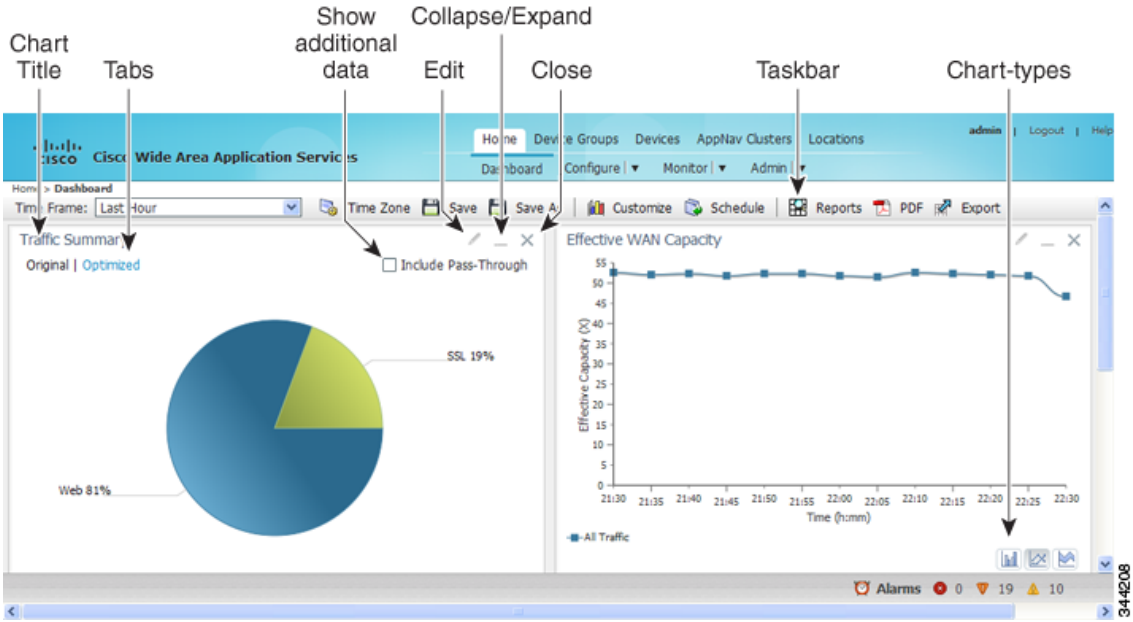

Taskbar icons and controls across the top of the dashboard or report allow you to do the following:

- Time Frame—Allows you to choose one of the several common time frames from the drop-down list:
  - **Last Hour**—Displays data for the past hour, in five-minute intervals (default). You can change the interval using the System.monitoring.collectRate configuration setting described in the "Modifying the Default System Configuration Properties" section on page 10-17.
  - **Last Day**—Displays data for the past day (in hourly intervals).
  - **Last Week**—Displays data for the past week (in daily intervals).
  - **Last Month**—Displays data for the past month (in daily intervals).
  - **Custom**—Enter starting and ending dates in the From and To fields. Click the calendar icon to choose dates from a popup calendar.

  The time frame setting is stored individually for each report and Central Manager user. Additionally, the System.monitoring.timeFrameSettings system property controls the system default time frame setting (see the "Modifying the Default System Configuration Properties" section on page 10-17).

Note If you create a chart with a custom date setting that spans more than two months back from the current date, the most recent two months of data are plotted with daily data and all previous months are plotted with aggregated monthly data. The chart might appear to have a large drop in traffic for the most recent two months because the daily traffic totals are likely to be much smaller than the monthly traffic totals; however, this difference is normal.

- Time Zone—Allows you to choose one of the following options from the Time Zone drop-down list:
  - **UTC**—Sets the time zone of the report to UTC.
  - **CM Local Time**—Sets the time zone of the report to the time zone of the WAAS Central Manager (default).

When you change the time zone, the change applies globally to all reports. The time zone setting is stored individually for each Central Manager user.

- Save—Saves the dashboard or report with its current settings. The next time you view it, it is displayed with these settings.
- Save As—Saves the report with its current settings under a new name. A popup window allows you to enter a report name and an optional description. You can enter only the following characters: numbers, letters, spaces, periods, hyphens, and underscores. The report will be available in the **Monitor > Reports > Reports Central** window.
- Customize—Allows you to add a chart or table to a dashboard or report. For information on adding a chart or table, see the "Adding a Chart or Table" section on page 17-13.
- Schedule—Allows you to schedule reports to be generated once or periodically such as hourly, daily, weekly, or monthly. When a scheduled report is generated, you can have a PDF copy of the report e-mailed to you automatically.
  - In the Date field, enter the schedule date in the format DD/MM/YYYY or click the calendar icon to display a calendar popup window from which to choose the date.
  - In the Hours drop-down list, choose the hours. The time represents the local time at the WAAS Central Manager.
  - In the Minutes drop-down list, choose the minutes. The time represents the local time at the WAAS Central Manager.
  - In the Frequency drop-down list, choose **Once**, **Hourly**, **Daily**, **Weekly**, or **Monthly** for the report frequency.
  - In the No. of Reports field, enter the number of times that a reoccurring report is to be generated. After a report is generated a specified number of times, the report is no longer generated.
  - In the Email Id(s) field, enter the e-mail addresses of the report recipients, separated by commas.
  - In the Email Subject field, enter the subject of the e-mail message.
- Reports—Allows you to view the scheduled reports. For instructions to view scheduled reports, see the "Managing Scheduled Reports" section on page 17-47.
- PDF—Generates a PDF format of the report, including the charts and table data.
- Export—Exports the chart statistical data to a CSV file. The statistical data shown in charts is rounded to whole units (KB, MB, or GB), while the exported data contains exact byte values.

Controls at the top of individual charts allow you to customize the chart as follows (not all controls are available in every chart):

- Chart title—Allows you to click and drag to move the chart to a different location in the report pane.
- Edit icon—Allows you to edit the chart settings as described in the "Configuring Chart Settings" section on page 17-14.
- Collapse/Expand icon—Allows you to collapse or expand the chart. When a chart is collapsed, this icon changes to Expand, which restores the chart to its normal size.
- Close icon—Closes the chart.
- Tabs—Allows you to have a choice of multiple tab views that you can access by clicking on the desired tab name (not all charts have this feature).
- Check box to show additional data—Allows you to check the box labeled with an optional data statistic to include the data in the chart (not all charts have this feature).

Chart-type icons at the bottom of individual charts allow you to choose the chart type as follows (not all controls are available in every chart): column chart, line chart, area chart, stacked line chart, stacked area chart.

Adding a Chart or Table

To add a chart or table to a dashboard or report, follow these steps:

Step 1 From the dashboard or report chart panel, click the **Customize** icon in the taskbar. The Customize window is displayed, as shown in Figure 17-7.

Figure 17-7 Customize Window

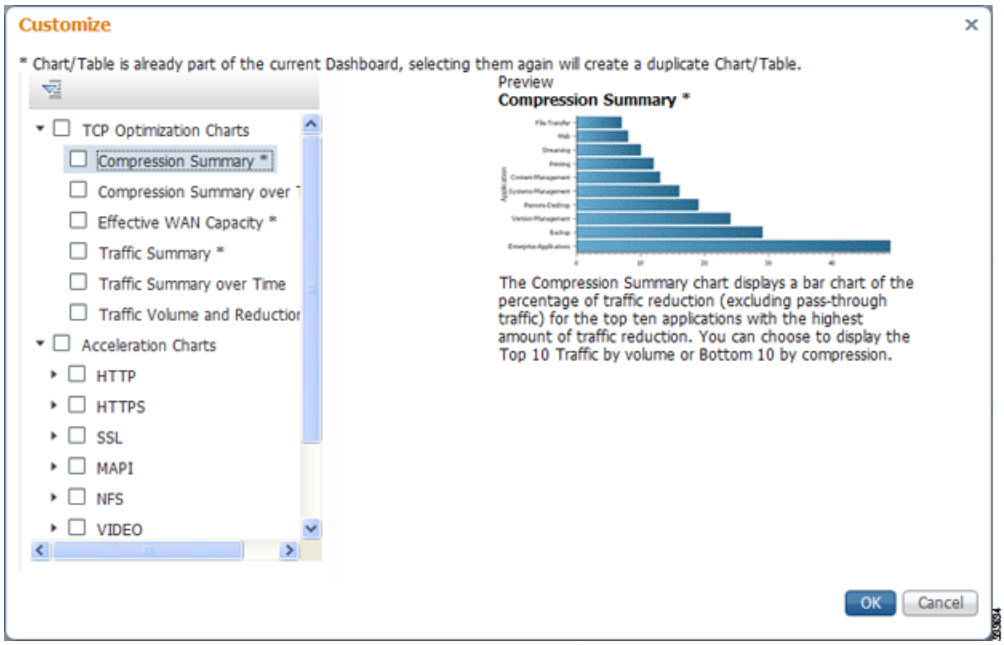

Step 2 Expand any of the chart categories by clicking on the small triangle next to the category.

Step 3 Check the box next to each chart or table that you want to display in the report. Individual charts are described in more detail in the "Chart and Table Descriptions" section on page 17-14.

Charts that are currently included in the dashboard or report are marked with an asterisk (*). A report can contain a maximum of eight charts and tables (the Network Summary report can contain 12).

Note At the WAAS Express device level, only charts for supported accelerators are available.

Step 4 To preview a chart, click on the chart title. The preview is displayed on the right of the pane.

Step 5 Click **OK**.

If you want to delete a chart or table from a dashboard or report, click the **Close** button on the chart and save the report.

Configuring Chart Settings

To configure the data presented in a chart, follow these steps:

Step 1 Click the **Edit** icon in the upper right corner of a chart. The Settings window is displayed.

Note Not all settings are available for all chart types.

Step 2 (Optional) From the Traffic Direction drop-down list, choose one of the following options:

- **Bidirectional**—Includes LAN to WAN traffic as well as WAN to LAN traffic traveling through this WAAS device.
- **Inbound**—Includes traffic from the WAN to the client through this WAAS device
- **Outbound**—Includes traffic traveling from a client to the WAN through this WAAS device.

Step 3 (Optional) In the Select Series For drop-down list, choose one of the following:

- **Application**—The chart data is based on application statistics.
- **Classifier**—The chart data is based on classifier (class map) statistics.

Step 4 (Optional) In the Application or Classifier list, check the box next to the applications or classifiers whose statistics you want to include in the chart data. To include all applications, check **All Traffic**. You can filter the list items by using the Quick Filter above the list. This list is available only for some chart types.

Step 5 (Optional) Some charts have other types of data series from which to choose. Check the box next to each data series that you want to include in the chart data.

Step 6 Click **OK**.

Note Data collection for applications and classifiers happens at slightly different times in the Central Manager, so the statistics can be different when viewing the same time period for an application and a classifier that report similar data.

Chart and Table Descriptions

This section describes the charts and tables that you can choose to include in a dashboard or report. The following categories are available:

- TCP Optimization Charts, page 17-15
- Acceleration Charts, page 17-16
- Connection Trend Charts, page 17-25
- AppNav Charts, page 17-26
- Platform Charts, page 17-27
- Statistics Details (Tables), page 17-28

All charts are plotted using the Central Manager local time zone, unless the chart settings were customized to use a different time zone.

Note At the device level for WAAS Express devices, only charts for supported accelerators are available. In all charts, pass-through traffic for WAAS Express devices is considered as zero.

TCP Optimization Charts

The following TCP optimization charts are available:

Compression Summary

The Compression Summary chart displays a bar chart of the percentage of traffic reduction (excluding pass-through traffic) for the top ten applications with the highest percentage of traffic reduction. Two additional tabs allow you to see the compression of the top ten applications by volume and the bottom ten applications with the lowest compression.

Formula:

% Reduction Excluding Pass-Through = (Original Excluding Pass-Through – Optimized) / (Original Excluding Pass-Through)

Compression Summary Over Time

The Compression Summary Over Time chart displays a graph of the percentage of total traffic that was reduced by using the WAAS optimization techniques. This chart excludes pass-through traffic in the results. You can customize the chart by choosing specific applications to include; the default is all traffic.

Formula:

% Reduction = (Original Excluding Pass-Through – Optimized) / (Original Excluding Pass-Through)

Effective WAN Capacity

The Effective WAN Capacity chart displays the effective increased bandwidth capacity of the WAN link as a result of WAAS optimization, as a value between 1X (times) and 100X. You can choose which applications to include; the default is all traffic.

Formula:

Effective WAN Capacity = 1 / (1-% Reduction Excluding Pass-Through)

% Reduction Excluding Pass-Through = (Original Excluding Pass-Through – Optimized) / (Original Excluding Pass-Through)

Throughput Summary

The Throughput Summary chart displays the amount of average and peak throughput for the LAN-to-WAN (outbound) or WAN-to-LAN (inbound) directions depending on the selected tab. The throughput units (kbps, mbps, or gbps) at the left side vary depending on the range. The Peak Throughput series is not applicable for Last Hour graphs. This chart is available only at the device and location levels.

Formula:

% Reduction Excluding Pass-Through = (Original Excluding Pass-Through – Optimized) / (Original Excluding Pass-Through)

Traffic Summary

The Traffic Summary chart displays the top nine applications that have the highest percentage of traffic as seen by WAAS. Each section in the pie chart represents an application as a percentage of the total traffic on your network or device. Unclassified, unmonitored, and applications with less than 2 percent of the total traffic are grouped together into a tenth category named Other Traffic (shown only if it totals at least 0.1 percent of all traffic). You can choose to display Original traffic or Optimized traffic by clicking the tab, and you can include pass-through traffic by checking the Include Pass-Through check box.

Formula:

(App Traffic/ Total Traffic) * 100

App Traffic is the Original traffic (Original Excluding Pass-Through) or Optimized traffic (Optimized Excluding Pass-Through) flowing for an application.

Traffic Summary Over Time

The Traffic Summary Over Time chart displays a graph of the amount of original or optimized traffic, depending on the selected tab, and you can include pass-through traffic by checking the Pass-Through check box. You can customize the chart by choosing specific applications to include; the default is all traffic.

Traffic Volume and Reduction

The Traffic Volume and Reduction chart compares the amount of original and optimized traffic in a bar chart and displays the percentage of traffic reduction as a line. Pass-through traffic is excluded. The traffic units (bytes, KB, MB, or GB) at the right side depend upon the range. The percentage of traffic reduction is shown at the left side of the chart. You can customize the chart by choosing specific applications to include; the default is all traffic.

Formula:

% Reduction Excluding Pass-Through = (Original Excluding Pass-Through – Optimized) / (Original Excluding Pass-Through)

Acceleration Charts

This section describes these charts:

HTTP Acceleration Charts

This section describes these charts:

HTTP: Connection Details

The HTTP Connection Details chart displays the HTTP session connection statistics, showing the average number of active HTTP connections per device (at the device level for the last hour it shows the exact number). Click the **Details** tab to display the newly handled HTTP connections, optimized connections, dropped connections, and handed off connections over time.

HTTP: Effective WAN Capacity

The HTTP Effective WAN Capacity chart displays the effective bandwidth capacity of the WAN link as a result of HTTP acceleration, as a multiplier of its base capacity. The capacity data for all traffic and HTTP traffic is shown.

Note If the chart has no data, monitoring may be disabled for the application definition that includes this type of traffic. Check that monitoring is enabled for the Web application.

HTTP: Estimated Time Savings

The HTTP Estimated Time Savings chart displays a graph of the estimated percentage of response time saved by the HTTP accelerator due to fast connection reuse and metadata caching.

HTTP: Optimization Count

The HTTP Optimization Count chart displays a graph of the number of different kinds of optimizations performed by the HTTP accelerator, which are displayed in different colors. The optimizations included in this chart are fast connection reuse and metadata caching.

HTTP: Optimization Techniques

The HTTP Optimization Techniques pie chart displays the different kinds of optimizations performed by the HTTP accelerator. The optimizations included in this chart are fast connection reuse, metadata caching, suppressed server compression, and DRE hinting.

HTTP: Response Time Savings

The HTTP Response Time Savings chart displays a graph of the round-trip response time saved by the HTTP accelerator due to metadata caching and fast connection reuse optimizations, which are displayed in different colors. The time units (milliseconds, seconds, or minutes) at the left side depend on the range.

HTTPS Acceleration Charts

This section describes the following charts:

HTTPS: Connection Details

The HTTPS Connection Details chart displays the HTTPS session connection statistics, showing the average number of active HTTPS connections per device (at the device level for the last hour it shows the exact number). Click the **Details** tab to display the newly handled HTTPS connections and optimized connections.

HTTPS: Effective WAN Capacity

The HTTPS Effective WAN Capacity chart displays the effective bandwidth capacity of the WAN link as a result of HTTP acceleration, as a multiplier of its base capacity. The capacity data for all traffic and SSL traffic (which includes HTTPS traffic) is shown.

Note If the chart has no data, monitoring may be disabled for the application definition that includes this type of traffic. Make sure that monitoring is enabled for the SSL application.

HTTPS: Estimated Time Savings

The HTTPS Estimated Time Savings chart displays the estimated percentage of response time saved by using metadata caching for HTTPS connections.

HTTPS: Optimization Count

The HTTPS Optimization Count chart displays a graph of the number of different kinds of metadata caching optimizations performed by the HTTPS accelerator, which are displayed in different colors.

HTTPS: Optimization Techniques

The HTTPS Optimization Techniques pie chart displays the different kinds of optimizations performed by the HTTPS accelerator. The optimizations included in this chart are metadata caching, suppressed server compression, and DRE hinting.

HTTPS: Response Time Savings

The HTTPS Response Time Savings chart displays a graph of the round-trip response time saved by the HTTPS accelerator due to metadata caching optimizations, which are displayed in different colors. The time units (milliseconds, seconds, or minutes) at the left side depend on the range.

Video Acceleration Charts

This section describes these charts:

- Video: Acceleration Bypass Reason, page 17-19
- Video: Connection Details, page 17-19
- Video: Effective WAN Capacity, page 17-19
- Video: Stream Optimization, page 17-20

Video: Acceleration Bypass Reason

The Video Acceleration Bypass Reason pie chart displays the reasons that video traffic is not accelerated: Windows Media VOD, aggregate bitrate overload, other reasons, stream bitrate overload, session count overload, and unsupported transmission type (which means an unsupported transport, unsupported player, or unsupported protocol).

Video: Connection Details

The Video Connection Details chart displays the video session connection statistics, showing the average number of active streaming video connections per device (at the device level for the last hour it shows the exact number). Click the **Details** tab to display the newly handled video connections, optimized connections, handed off connections, and dropped connections.

Video: Effective WAN Capacity

The Video Effective WAN Capacity chart displays the effective bandwidth capacity of the WAN link as a result of video acceleration, as a multiplier of its base capacity. The capacity data for all traffic and streaming video traffic is shown.

Note If the chart has no data, monitoring may be disabled for the application definition that includes this type of traffic. Check that monitoring is enabled for the Streaming application.

Video: Stream Optimization

The Video Stream Optimization chart compares the amounts of traffic incoming from the WAN and outgoing to the LAN. The traffic units (bytes, KB, MB, or GB) at the left side depend on the range.

SSL Acceleration Charts

This section describes these charts:

- SSL: Acceleration Bypass Reason, page 17-20
- SSL: Connection Details, page 17-20
- SSL: Effective WAN Capacity, page 17-20

SSL: Acceleration Bypass Reason

The SSL Acceleration Bypass Reason pie chart displays the reasons that SSL traffic is not accelerated: version mismatch, unknown, nonmatching domain, cipher mismatch, revocation failure, certificate verification failure, other failure, and non-SSL traffic.

SSL: Connection Details

The SSL Connection Details chart displays the SSL session connection statistics, showing the average number of active SSL connections per device (at the device level for the last hour it shows the exact number). Click the **Details** tab to display the newly handled SSL connections, optimized connections, handed off connections, dropped connections, and HTTPS connections.

SSL: Effective WAN Capacity

The SSL Effective WAN Capacity chart displays the effective bandwidth capacity of the WAN link as a result of SSL acceleration, as a multiplier of its base capacity. The capacity data for all traffic and SSL traffic is shown.

Note If the chart has no data, monitoring may be disabled for the application definition that includes this type of traffic. Check that monitoring is enabled for the SSL application.

MAPI Acceleration Charts

This section describes these charts:

- MAPI: Acceleration Bypass Reason, page 17-21
- MAPI: Average Response Time Saved, page 17-21
- MAPI: Connection Details, page 17-21
- MAPI: Effective WAN Capacity, page 17-21
- MAPI: Request Optimization, page 17-21
- MAPI: Response Time Optimization, page 17-21
- MAPI: Current Accelerated Client Sessions, page 17-22

MAPI: Acceleration Bypass Reason

The MAPI Acceleration Bypass Reason pie chart displays the reasons that encrypted MAPI traffic is not accelerated: acceleration disabled, secret retriever disabled, unsupported cipher, unsupported authentication mechanism, misconfigured domain identity, failure in secret retrieval, general security failure, insufficient system resources, and recovery mode connections.

Click the **Non-Encrypted** tab to display the bypass reasons for unencrypted MAPI traffic: reservation failure (non-overload), reservation failure (overload), signed MAPI request, malformed RPC packet, handover request from peer, unsupported server version, user in denied list, unsupported client version, secured connections (encrypted), unsupported DCERPC protocol version, association group not tracked, and other.

MAPI: Average Response Time Saved

The MAPI Average Response Time Saved chart displays a graph of the estimated percentage of response time saved by the MAPI accelerator. The time units (microseconds, milliseconds, seconds, or minutes) at the left side depend upon the range.

MAPI: Connection Details

The MAPI Connection Details chart displays the MAPI session connection statistics, showing the average number of active MAPI connections per device (at the device level for the last hour it shows the exact number). Click the **Details** tab to display the newly handled MAPI connections, optimized connections, handed-off connections, and dropped connections. Click the **Optimized Encrypted vs Non-Encrypted** tab to display the new encrypted and unencrypted MAPI connections.

MAPI: Effective WAN Capacity

The MAPI Effective WAN Capacity chart displays the effective bandwidth capacity of the WAN link as a result of MAPI acceleration, as a multiplier of its base capacity. The capacity data for all traffic and MAPI traffic is shown.

Note If the chart has no data, monitoring may be disabled for the application definition that includes this type of traffic. Check that monitoring is enabled for the Email-and-Messaging application.

MAPI: Request Optimization

The MAPI Request Optimization chart displays the percentage of local and remote MAPI command responses. A local response is a response that is sent to the client from the local WAE. A remote response comes from the remote server. Click the **Encrypted vs Non-Encrypted** tab to display the percentage of local and remote responses for encrypted and unencrypted MAPI connections.

MAPI: Response Time Optimization

The MAPI Response Time Optimization chart compares the average time used for local and remote MAPI responses. The time units (microseconds, milliseconds, seconds, or minutes) at the left side depend upon the range. Click the **Encrypted vs Non-Encrypted** tab to display the average time used for local and remote responses for encrypted and unencrypted MAPI connections.

MAPI: Current Accelerated Client Sessions

The MAPI Current Accelerated Client Sessions pie chart displays the number of encrypted sessions currently being accelerated from different versions (2000, 2003, 2007, and 2010) of the Microsoft Outlook client. Click the **Non-Encrypted** tab to display the unencrypted session counts.

NFS Acceleration Charts

This section describes these charts:

NFS: Acceleration Bypass Reason

The NFS Acceleration Bypass Reason pie chart displays the reasons that NFS traffic is not accelerated: unknown authentication flavor or unknown NFS version.

NFS: Connection Details

The NFS Connection Details chart displays the NFS session connection statistics, showing the average number of active NFS connections per device (at the device level for the last hour it shows the exact number). Click the **Details** tab to display the newly handled NFS connections, optimized connections, handed-off connections, and dropped connections.

NFS: Effective WAN Capacity

The NFS Effective WAN Capacity chart displays the effective bandwidth capacity of the WAN link as a result of NFS acceleration, as a multiplier of its base capacity. The capacity data for all traffic and NFS traffic is shown.

Note If the chart has no data, monitoring may be disabled for the application definition that includes this type of traffic. Check that monitoring is enabled for the File-System application.

NFS: Estimated Time Savings

The NFS Estimated Time Savings chart displays a graph of the estimated percentage of response time saved by the NFS accelerator.

NFS: Request Optimization

The NFS Request Optimization chart displays the percentage of local and remote NFS command responses. A local response is a response that is sent to the client from the local WAE. A remote response comes from the remote server.

NFS: Response Time Optimization

The NFS Response Time Optimization chart compares the average time used for local and remote NFS responses. The time units (milliseconds, seconds, or minutes) at the left side depend upon the range.

NFS: Versions Detected

The NFS Versions Detected pie chart displays the number of NFS messages detected for each NFS version (2, 3, and 4). The NFS accelerator works with NFS version 3 traffic, so you will want to see this type of traffic for best results.

SMB Acceleration Charts

This section describes these charts:

SMB: Average Response Time Saved

The SMB Average Response Time Saved chart displays the average response time saved for SMB responses. The time units (milliseconds, seconds, or minutes) at the left side depend upon the range.

SMB: Client Average Throughput

The SMB Client Average Throughput chart displays the average client throughput for the SMB accelerator.

SMB: Connection Details

The SMB Connection Details chart displays the SMB session connection statistics, showing the average number of active SMB connections per device (at the device level for the last hour it shows the exact number). Click the **Details** tab to display the newly handled SMB connections, optimized connections, handed-off connections, dropped connections, and signed connections.

SMB: Effective WAN Capacity

The SMB Effective WAN Capacity chart displays the effective bandwidth capacity of the WAN link as a result of SMB acceleration, as a multiplier of its base capacity. The capacity data for all traffic and SMB traffic is shown.

Note If the chart has no data, monitoring may be disabled for the application definition that includes this type of traffic. Check that monitoring is enabled for the CIFS application.

SMB: Request Optimization

The SMB Request Optimization chart displays the percentage of SMB command responses that use the following optimizations: read ahead, metadata, write, and other.

SMB: Response Time Savings

The SMB Response Time Savings chart displays a graph of the round-trip response time saved by the SMB accelerator due to the following optimizations, which are displayed in different colors: read ahead, metadata, Microsoft Office, async write, named pipe, and other. The time units (milliseconds, seconds, or minutes) at the left side depend on the range.

SMB: Versions Detected

The SMB Versions Detected pie chart displays the number of SMB messages detected for each SMB version: SMB v1.0 optimized, SMB v1.0 unoptimized, SMB v1.0 signed, SMB v2.0 optimized, SMB v2.0 unoptimized, SMB v2.0 signed, SMB v2.1 optimized, SMB v2.1 unoptimized, and SMB v2.1 signed.

ICA Acceleration Charts

This section describes these charts:

- ICA: Client Versions, page 17-24
- ICA: Connection Details, page 17-24
- ICA: Effective WAN Capacity, page 17-24
- ICA: Unaccelerated Reasons, page 17-25

ICA: Client Versions

The ICA Client Versions pie chart displays the number of ICA messages detected for each ICA version: online plugin 11.0, online plugin 11.2, online plugin 12.0, online plugin 12.1, Citrix Receiver 13.0, and other.

ICA: Connection Details

The ICA Connection Details chart displays the ICA session connection statistics, showing the average number of active ICA connections per device (at the device level for the last hour it shows the exact number). Click the **Details** tab to display the newly handled ICA connections, optimized connections, handed-off connections, and dropped connections.

ICA: Effective WAN Capacity

The ICA Effective WAN Capacity chart displays the effective bandwidth capacity of the WAN link as a result of ICA acceleration, as a multiplier of its base capacity. The capacity data for all traffic and ICA traffic is shown.

Note If the chart has no data, monitoring may be disabled for the application definition that includes this type of traffic. Check that monitoring is enabled for the Citrix application.

ICA: Unaccelerated Reasons

The ICA Unaccelerated Reasons chart displays the reasons that ICA traffic is bypassed: unrecognized protocol, unsupported client version, CGP session ID unknown, client on denied list, no resource, and other. Click the **Dropped** tab to display the reasons that ICA traffic is dropped: unsupported client version, I/O error, no resource, AO parsing error, maximum sessions reached, and other.

CIFS Acceleration Charts for WAAS Express

This section describes these charts that are available only on WAAS Express devices:

- CIFS: Client Average Throughput, page 17-25
- CIFS: Connection Details, page 17-25
- CIFS: Effective WAN Capacity, page 17-25
- CIFS: Request Optimization, page 17-25

CIFS: Client Average Throughput

The CIFS Client Average Throughput chart displays the average client throughput for the WAAS Express CIFS accelerator.

CIFS: Connection Details

The CIFS Connection Details chart displays the CIFS session connection statistics for the WAAS Express CIFS accelerator, showing the average number of active CIFS connections per device (at the device level for the last hour it shows the exact number). Click the **Details** tab to display the newly handled CIFS connections, optimized connections, handed-off connections, and dropped connections.

CIFS: Effective WAN Capacity

The CIFS Effective WAN Capacity chart displays the effective bandwidth capacity of the WAN link as a result of WAAS Express CIFS acceleration, as a multiplier of its base capacity. The capacity data for all traffic and CIFS traffic is shown.

CIFS: Request Optimization

The CIFS Request Optimization chart displays the percentage of WAAS Express CIFS accelerator command responses that use the following optimizations: read ahead, metadata, write, and other.

Connection Trend Charts

This section describes these charts:

- Optimized Connections Over Time, page 17-26
- Optimized vs Pass-Through Connections, page 17-26

Optimized Connections Over Time

The Optimized Connections Over Time chart displays the number of optimized connections over the selected time period. You can show the number of MAPI reserved connections by checking the **MAPI Reserved Connections** check box. You can customize the chart by choosing specific applications to include; the default is all traffic.

This chart is available only when a specific WAAS device is selected and can be added only to the Connection Trend report.

Optimized vs Pass-Through Connections

The Optimized vs Pass-Through Connections chart displays the total number of optimized and pass-through connections on a device or on all devices in a location. You can show the device connection limit, which is the maximum number of connections a device can support, by checking the **Device Connection Limit** check box. This option is available only at the device level.

This chart is available only when a specific WAAS device or location is selected and can be added only to the Connection Trend report.

Formula:

Pass-Through Connections for a Device = Total Pass-Through Connections for all applications

Optimized Connections for a Device = Total Optimized Connections for all applications

AppNav Charts

This section describes these charts:

Total AppNav Traffic

The Total AppNav Traffic chart displays the total amount of distributed and pass-through traffic processed by the AppNav Cluster or ANC device. The units at the left side depend upon the range.

AppNav Policies

The AppNav Policies chart displays a graph of the amount of intercepted, distributed, or pass-through traffic processed by the AppNav Cluster or ANC device for each policy rule, depending on which tab you select. The units at the left side depend upon the range.

From the Show Details For drop-down list, you can select a policy rule for which to show details.

Top 10 AppNav Policies

The Top 10 AppNav Policies pie chart displays the amount of intercepted, distributed, or pass-through traffic processed by the AppNav Cluster or ANC device for the top nine policy rules with the most traffic, depending on which tab you select. Traffic for all other policy rules is grouped together into a tenth category named Other Traffic (shown only if it totals at least 0.1 percent of all traffic).

From the Show Details For drop-down list, you can select a policy rule for which to show details.

Top 10 WAAS Node Group Distribution

The Top 10 WAAS Node Group Distribution pie chart displays the top nine WNGs to which traffic is distributed. Traffic for all other WNGs is grouped together into a tenth category named Other Traffic (shown only if it totals at least 0.1 percent of all traffic).

From the Show Details For drop-down list, you can select a WNG for which to show details of the individual WNs.

WAAS Node Group Distribution

The WAAS Node Group Distribution chart displays a graph of the amount of traffic distributed to each of the WAAS node groups (WNGs). The units at the left side depend upon the range.

From the Show Details For drop-down list, you can select a WNG for which to show details of the individual WAAS nodes (WNs).

Pass-Through Reasons

The Pass-Through Reasons chart displays a graph of the amount of pass-through traffic for each of the pass-through reasons. The units at the left side depend upon the range.

From the Show Details For drop-down list, you can select a reason for which to show details.

Top 10 Pass-Through Reasons

The Top 10 Pass-Through Reasons pie chart displays the top nine reasons that traffic is passed through. Traffic for all other reasons is grouped together into a tenth category named Other Traffic (shown only if it totals at least 0.1 percent of all traffic).

From the Show Details For drop-down list, you can select a reason for which to show details.

Platform Charts

This section describes these charts:

- CPU Utilization, page 17-28
- Disk Utilization, page 17-28

CPU Utilization

The CPU Utilization chart displays the percentage of CPU utilization for the device. This chart is available only when a specific WAAS device is selected. This chart can be added only to the **Monitor > Reports > Reports Central > Resource Utilization** report page.

Disk Utilization

The Disk Utilization chart displays the percentage of disk utilization for the device. This chart is available only when a specific WAAS device is selected. This chart can be added only to the **Monitor > Reports > Reports Central > Resource Utilization** report page.

Statistics Details (Tables)

The following statistics details tables are available:

- Traffic Summary Table, page 17-29
- Network Application Traffic Details Table, page 17-30
- HTTP Acceleration Statistics Table, page 17-30
- HTTPS Acceleration Statistics Table, page 17-30
- ICA Acceleration Statistics Table, page 17-31
- MAPI Acceleration Statistics Table, page 17-31
- NFS Acceleration Statistics Table, page 17-32
- SMB Acceleration Statistics Table, page 17-33
- SSL Acceleration Statistics Table, page 17-33
- Video Acceleration Statistics Table, page 17-34
- CIFS Acceleration Statistics Table for WAAS Express, page 17-34

You can sort the tables by clicking on any column heading to sort on data in that column. A small triangle control appears in the heading to indicate that the table is sorted on this column. Click the triangle to reverse the sort order for the column.

For some values, different formulas are used at the system and device levels, and these formulas are noted in the table descriptions. The terms used in the tables are defined as follows:

- Original Inbound—Traffic that is entering the WAAS device from the LAN (clients) and needs to be optimized before being sent out on the WAN to a peer WAAS device.
- Original Outbound—Traffic that is exiting the WAAS device to the LAN (clients) after being received on the WAN from a peer WAAS device.
- Optimized Inbound—Traffic that is entering the WAAS device from the WAN and needs to be processed (deoptimized) before being sent out on the LAN to clients.
- Optimized Outbound—Traffic that is exiting the WAAS device to the WAN and a peer WAAS device after being optimized.
- Pass-Through—Traffic that is being passed through the WAAS device and not optimized.

To get the statistics at the system, location, and device group levels, the Original Inbound, Original Outbound, Optimized Inbound, Optimized Outbound, Pass-through Client, and Pass-through Server bytes of all devices are added together. The Reduction % and Effective Capacity values are calculated using these added values of all devices.

Traffic Summary Table

This table is called the Network Traffic Summary, Device Traffic Summary, or Location Traffic Summary, depending on the context and it displays a summary of traffic.

At the system and location levels, each row in the table displays the total traffic information for each device that is registered to this Central Manager or is in this location. At the device level, each row in the table displays the total traffic information for each application defined on the device. The data is described in Table 17-5.

Table 17-5 Traffic Summary Table

Table Column	Description and Formulas Used to Calculate Value
Device	The device name. (Appears only at the system and location levels.)
Application	The application name. (Appears only at the device level. At the system level, use the Network Application Traffic Details Table to get this information.)
Original Traffic (Excludes Pass-Through)	Reports the amount of original traffic, excluding pass-through traffic. System: (Original Outbound + Original Inbound)/2 Device/Device Group: Original Inbound + Original Outbound
Optimized Traffic (Excludes Pass-Through)	Reports the amount of optimized traffic, excluding pass-through traffic. System: (Optimized Inbound + Optimized Outbound)/2 Device/Device Group: Optimized Outbound + Optimized Inbound
Pass-Through Traffic	Reports the amount of pass-through traffic. This value is not applicable for WAAS Express devices. System: (Pass-through Client + Pass-through Server)/2 Device/Device Group: Pass-through Client + Pass-through Server An asterisk (*) in the column heading indicates that a device whose data is included in this table is configured as a serial peer with another device and optimization is disabled between those two peer devices. The amount of pass-through traffic shown may be more than what is expected because the device passes through traffic coming from its peer (for more information, see the "Information About Clustering Inline WAEs" section on page 5-53).[1]
Reduction (%)	Reports the percentage of bytes saved, considering only optimized traffic. (Original Excl Pass-through – (Optimized)) * 100 / (Original Excl Pass-through)
Effective Capacity	Reports the effective bandwidth capacity of the WAN link as a result of optimization, as a multiplier of its base capacity, considering only optimized traffic. 1/(1 – % Reduction Excl Pass-through)

1. The number in the Pass-Through Traffic column represents the amount of traffic that is passed through that particular WAE (or for a location report, all the devices in the location). If the device is part of a serial inline cluster (that is, configured as a non-optimizing peer with another device), the traffic that is shown as pass-through on one device may have been optimized by another device in the serial cluster. It is useful to know the amount of traffic that is not optimized by either of the devices in the cluster (in other words, passed through the entire cluster).

   When the device closer to the LAN is not overloaded, the pass through numbers on that device accurately represent the overall pass-through traffic. But, if that device goes into overload, the second device in the cluster starts optimizing traffic that was passed through by the first one, which needs to be accounted for. In this case, the overall pass-through numbers for the cluster can be obtained as follows. Note that this calculation has to be done even if the first device went into overload in the past and came out of it.

   Consider that W1 and W2 are part of a serial cluster and W1 is toward the LAN (closer to the client if the cluster is in the branch, or closer to the server if the cluster is in the data center) and W2 is toward the WAN. The amount of traffic that is passed through the cluster without optimization by either W1 or W2 can be obtained by the following formula: (W1 pass-through traffic) – (W2 original traffic)

Network Application Traffic Details Table

The Network Application Traffic Details table is available at the system level and displays the total traffic information for each application. The data is the same as described in Table 17-5 (except there is no Device column in this table).

HTTP Acceleration Statistics Table

The HTTP Acceleration Statistics table is available at the system and device levels and displays HTTP acceleration details. The data is described in Table 17-6.

Table 17-6 HTTP Acceleration Statistics Table

Table Column	Description and Formulas Used to Calculate Value
Device	The device name. (Appears only at the system level.)
Start Time and End Time	Start and end times for the time period. (Appears only at the device level.)
New Connections Handled	Reports the number of HTTP connections handled for the time period.
Average Active Connections/ Active Connections	Reports the average active number of connections currently being handled by the HTTP accelerator at the system level. At other levels, reports the number of active connections.
New Bypassed Connections	Reports the number of connections initially received by the HTTP accelerator and then pushed down to the generic accelerator.
Total Time Saved	Reports the amount of time saved due to HTTP optimization.
Total Round-Trip Time	Reports the total round-trip time for all connections plus the time for remotely served metadata cache misses.
% Time Saved	Reports the percentage of connection time saved for all aggregated samples. Total Time Saved / (Total Time Saved + Total Round Trip Time For All Connections + Total time for all remotely served metadata cache misses)

HTTPS Acceleration Statistics Table

The HTTPS Acceleration Statistics table is available at the system and device levels and displays HTTPS acceleration details. The data is described in Table 17-7.

Table 17-7 HTTPS Acceleration Statistics Table

Table Column	Description and Formulas Used to Calculate Value
Device	The device name. (Appears only at the system level.)
Start Time and End Time	Start and end times for the time period. (Appears only at the device level.)
New Connections Handled	Reports the number of HTTPS connections handled for the time period.
Average Active Connections/ Active Connections	Reports the average number of connections currently being handled by the HTTP/SSL accelerator at the system level. At other levels, reports the number of active connections.
Total Time Saved	Reports the amount of time saved due to HTTPS optimization.
Total Round-Trip Time	Reports the total round-trip time for all connections plus the time for remotely served metadata cache misses.
% Time Saved	Reports the percentage of connection time saved for all aggregated samples. Total Time Saved by cache hits / (Total Time Saved by cache hits + Total time for all remotely served metadata cache misses)

ICA Acceleration Statistics Table

The ICA Acceleration Statistics table is available at the system and device levels and displays ICA acceleration details. The data is described in Table 17-8.

Table 17-8 ICA Acceleration Statistics Table

Table Column	Description and Formulas Used to Calculate Value
Device	The device name. (Appears only at the system level. WAAS Express devices are not included.)
Start Time and End Time	Start and end times for the time period. (Appears only at the device level.)
New Connections Handled	Reports the number of ICA connections handled for the time period.
Average Active Connections/ Active Connections	Reports the average number of connections currently being handled by the ICA accelerator at the system level. At other levels, reports the number of active connections.
Dropped Connections	Reports the number of connections dropped by the ICA accelerator.
Bypassed Connections	Reports the number of connections initially received by the ICA accelerator and then pushed down to the generic accelerator.

MAPI Acceleration Statistics Table

The MAPI Acceleration Statistics table is available at the system and device levels and displays MAPI acceleration details. The data is described in Table 17-9.

Table 17-9 MAPI Acceleration Statistics Table

Table Column	Description and Formulas Used to Calculate Value
Device	The device name. (Appears only at the system level. WAAS Express devices are not included.)
Start Time and End Time	Start and end times for the time period. (Appears only at the device level.)
New Connections Handled	Reports the number of MAPI connections handled for the time period.
Average Active Connections/ Active Connections	Reports the average number of connections currently being handled by the MAPI accelerator at the system level. At other levels, reports the number of active connections.
New Bypassed Connections	Reports the number of connections initially received by the MAPI accelerator and then pushed down to the generic accelerator.
New Local Request Count	Reports the number of client requests handled locally by the WAE.
Avg. Local Response Time	Reports the average time used for local responses, in microseconds.
New Remote Request Count	Reports the number of client requests handled remotely over the WAN.
Avg. Remote Response Time	Reports the average time used for remote responses, in microseconds.
Average Time Saved	Reports the average connection time saved for all aggregated samples, in microseconds.

NFS Acceleration Statistics Table

The NFS Acceleration Statistics table is available at the system and device levels and displays NFS acceleration details. The data is described in Table 17-10.

Table 17-10 NFS Acceleration Statistics Table

Table Column	Description and Formulas Used to Calculate Value
Device	The device name. (Appears only at the system level. WAAS Express devices are not included.)
Start Time and End Time	Start and end times for the time period. (Appears only at the device level.)
New Connections Handled	Reports the number of NFS connections handled for the time period.
Average Active Connections/ Active Connections	Reports the average number of connections currently being handled by the NFS accelerator at the system level. At other levels, reports the number of active connections.
New Bypassed Connections	Reports the number of connections initially received by the NFS accelerator and then pushed down to the generic accelerator.
New Local Request Count	Reports the number of client requests handled locally by the WAE.
Avg. Local Response Time	Reports the average time used for local responses, in milliseconds.
New Remote Request Count	Reports the number of client requests handled remotely over the WAN.

Table 17-10 NFS Acceleration Statistics Table

Table Column	Description and Formulas Used to Calculate Value
Avg. Remote Response Time	Reports the average time used for remote responses, in milliseconds.
% Time Saved	Reports the percentage of connection time saved for all aggregated samples. (Down – Up) * 100 / (Down) If(Down != 0) where: Down = (New local request count + New remote request count) * Avg. local response time Up = ((New local request count * Avg. local response time) + (New remote request count * Avg. remote response time))

SMB Acceleration Statistics Table

The SMB Acceleration Statistics table is available at the system and device levels and displays SMB acceleration details. The data is described in Table 17-11.

Table 17-11 SMB Acceleration Statistics Table

Table Column	Description and Formulas Used to Calculate Value
Device	The device name. (Appears only at the system level. WAAS Express devices are not included.)
Start Time and End Time	Start and end times for the time period. (Appears only at the device level.)
New Connections Handled	Reports the number of SMB connections handled for the time period.
Average Active Connections/ Active Connections	Reports the average number of connections currently being handled by the SMB accelerator at the system level. At other levels, reports the number of active connections.
Bypassed Connections	Reports the number of connections initially received by the SMB accelerator and then pushed down to the generic accelerator.
Total Time Saved	Reports the amount of time saved due to SMB optimization.
Total Round-Trip Time	Reports the total round-trip time for all connections plus the time for remotely served metadata cache misses.
% Time Saved	Reports the percentage of connection time saved for all aggregated samples. Total Time Saved by cache hits / (Total Time Saved by cache hits + Total Time for all remotely served metadata cache misses)

SSL Acceleration Statistics Table

The SSL Acceleration Statistics table is available at the system and device levels and displays SSL acceleration details. The data is described in Table 17-12.

Table 17-12 SSL Acceleration Statistics Table

Table Column	Description
Device	The device name. (Appears only at the system level.)
Start Time and End Time	Start and end times for the time period. (Appears only at the device level.)
New Connections Handled	Reports the number of SSL connections handled for the time period.
Average Active Connections/ Active Connections	Reports the average number of connections currently being handled by the SSL accelerator at the system level. At other levels, reports the number of active connections.
New HTTPS Connections Handled	Reports the number of HTTPS connections handled by the SSL accelerator.
Dropped Connections	Reports the number of connections dropped by the SSL accelerator.
Bypassed Connections	Reports the number of connections initially received by the SSL accelerator and then pushed down to the generic accelerator.

Video Acceleration Statistics Table

The Video Acceleration Statistics table is available at the system and device levels and displays video acceleration details. The data is described in Table 17-13.

Table 17-13 Video Acceleration Statistics Table

Table Column	Description
Device	The device name. (Appears only at the system level.)
Start Time and End Time	Start and end times for the time period. (Appears only at the device level.)
New Connections Handled	Reports the number of video connections handled for the time period.
Average Active Connections/ Active Connections	Reports the average number of connections currently being handled by the video accelerator at the system level. At other levels, reports the number of active connections.
New Bypassed Connections	Reports the number of connections initially received by the video accelerator and then pushed down to the generic accelerator.

CIFS Acceleration Statistics Table for WAAS Express

The CIFS Acceleration Statistics table displays CIFS acceleration details for a WAAS Express device. The data is described in Table 17-14.

Table 17-14 CIFS Acceleration Statistics Table

Table Column	Description and Formulas Used to Calculate Value
Start Time and End Time	Start and end times for the time period. (Appears only at the device level.)
New Connections Handled	Reports the number of CIFS connections handled for the time period.
Average Active Connections/ Active Connections	Reports the average number of connections currently being handled by the CIFS accelerator at the system level. At other levels, reports the number of active connections.

Table 17-14 CIFS Acceleration Statistics Table

Table Column	Description and Formulas Used to Calculate Value
Bypassed Connections	Reports the number of connections initially received by the CIFS accelerator and then pushed down to the generic accelerator.
Total Time Saved	Reports the amount of time saved due to CIFS optimization.
Total Round-Trip Time	Reports the total round-trip time for all connections plus the time for remotely served metadata cache misses.
% Time Saved	Reports the percentage of connection time saved for all aggregated samples. Total Time Saved by cache hits / (Total Time Saved by cache hits + Total time for all remotely served metadata cache misses)

Using Predefined Reports to Monitor WAAS

The WAAS Central Manager includes a number of predefined reports that you can use to monitor the system operation. These reports are available in the Monitor menu. The reports consist of a combination of specific charts and graphs and a statistical table displayed in the lower part of the window.

You can customize these predefined reports by editing them with the Manage Report function available in the Monitor menu, as described in the "Viewing and Editing Reports" section on page 17-45.

The following predefined reports are available at the WAAS system level, the AppNav Cluster level, the location level, and the device level:

- Optimization
  - TCP Summary Report, page 17-36
- Acceleration (not all are available at the WAAS Express device level)
  - HTTP Acceleration Report, page 17-37
  - HTTPS Acceleration Report, page 17-37
  - Video Acceleration Report, page 17-37
  - SSL Acceleration Report, page 17-38
  - MAPI Acceleration Report, page 17-38
  - NFS Acceleration Report, page 17-38
  - SMB Acceleration Report, page 17-38
  - ICA Acceleration Report, page 17-39

The following predefined report is available only at the WAAS System level:

- Network > Summary Report, page 17-39

The following predefined report is available only at the WAAS System level and the device level:

- Network/Peers > Topology Report, page 17-40

The following predefined report is available only at the device level and the location level:

- Optimization > Connection Trend Report, page 17-40

The following predefined reports are available only at the device level:

- Optimization
  - Connections Statistics Report, page 17-40
- Acceleration
  - CIFS Acceleration Report, page 17-41 (not available for a WAAS Express device)
  - CIFS Acceleration Report for WAAS Express, page 17-42 (available only for a WAAS Express device)
- Platform (not available at the WAAS Express device level)
  - Resource Utilization Report, page 17-42
  - Disks Report, page 17-42

The following predefined reports are available only at the AppNav Cluster level and at the device level for AppNav Controller devices:

- AppNav > AppNav Report, page 17-43

Note In a WAAS network where there are 1000 or more WAEs, there may be a delay of up to 90 seconds to redisplay the table when you click a table column to resort any system level report table. You may experience a similar delay when you click the Print icon in the taskbar, before you see the report.

Location Level Reports

Location level reports aggregate data from all the WAEs present in a particular location. For more information about locations, see the "Working with Device Locations" section on page 3-9.

To view a location level report, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Locations >** *location-name*.

Step 2 Choose **Monitor** and choose the report from the Optimization or Acceleration categories.

When scheduling any report, you can also select one or more locations and the report will include data from all devices within the selected locations. For more information, see the "Scheduling Reports" section on page 17-46.

The maximum number of devices supported in a location level report is 25 by default. This number is configurable up to 250 by the System.monitoring.maxDevicePerLocation system property. For more information, see the "Modifying the Default System Configuration Properties" section on page 10-17.

TCP Summary Report

The TCP Summary report displays a summary of all traffic. The following charts and tables are included:

- Traffic Summary, page 17-16
- Effective WAN Capacity, page 17-15
- Traffic Volume and Reduction, page 17-16

- Compression Summary, page 17-15
- Traffic Summary Over Time, page 17-16
- Compression Summary Over Time, page 17-15
- Throughput Summary, page 17-16 (included only at the device and location levels)
- Traffic Summary Table, page 17-29

HTTP Acceleration Report

The HTTP Acceleration report displays the HTTP acceleration statistics. The following charts and tables are included:

- HTTP: Estimated Time Savings, page 17-17
- HTTP: Effective WAN Capacity, page 17-17
- HTTP: Connection Details, page 17-17
- HTTP: Response Time Savings, page 17-18
- HTTP: Optimization Count, page 17-18
- HTTP: Optimization Techniques, page 17-18
- HTTP Acceleration Statistics Table, page 17-30

HTTPS Acceleration Report

The HTTPS Acceleration report displays the HTTPS acceleration statistics. The following charts and tables are included:

- HTTPS: Estimated Time Savings, page 17-18
- HTTPS: Effective WAN Capacity, page 17-18
- HTTPS: Connection Details, page 17-18
- HTTPS: Response Time Savings, page 17-19
- HTTPS: Optimization Count, page 17-19
- HTTPS: Optimization Techniques, page 17-19
- HTTPS Acceleration Statistics Table, page 17-30

Video Acceleration Report

The Video Acceleration report displays the video acceleration statistics. The following charts and tables are included:

- Video: Stream Optimization, page 17-20
- Video: Effective WAN Capacity, page 17-19
- Video: Connection Details, page 17-19
- Video: Acceleration Bypass Reason, page 17-19
- Video Acceleration Statistics Table, page 17-34

SSL Acceleration Report

The SSL Acceleration report displays the SSL acceleration statistics. The following charts and tables are included:

- SSL: Connection Details, page 17-20
- SSL: Effective WAN Capacity, page 17-20
- SSL: Acceleration Bypass Reason, page 17-20
- SSL Acceleration Statistics Table, page 17-33

MAPI Acceleration Report

The MAPI Acceleration report displays the MAPI acceleration statistics. The following charts and tables are included:

- MAPI: Average Response Time Saved, page 17-21
- MAPI: Effective WAN Capacity, page 17-21
- MAPI: Connection Details, page 17-21
- MAPI: Request Optimization, page 17-21
- MAPI: Response Time Optimization, page 17-21
- MAPI: Current Accelerated Client Sessions, page 17-22
- MAPI: Acceleration Bypass Reason, page 17-21
- MAPI Acceleration Statistics Table, page 17-31

NFS Acceleration Report

The NFS Acceleration report displays the NFS acceleration statistics. The following charts and tables are included:

- NFS: Estimated Time Savings, page 17-22
- NFS: Effective WAN Capacity, page 17-22
- NFS: Connection Details, page 17-22
- NFS: Request Optimization, page 17-22
- NFS: Response Time Optimization, page 17-23
- NFS: Versions Detected, page 17-23
- NFS: Acceleration Bypass Reason, page 17-22
- NFS Acceleration Statistics Table, page 17-32

SMB Acceleration Report

The SMB Acceleration report displays the SMB acceleration statistics. The following charts and tables are included:

- SMB: Average Response Time Saved, page 17-23

- SMB: Effective WAN Capacity, page 17-23
- SMB: Connection Details, page 17-23
- SMB: Request Optimization, page 17-24
- SMB: Response Time Savings, page 17-24
- SMB: Client Average Throughput, page 17-23
- SMB: Versions Detected, page 17-24
- SMB Acceleration Statistics Table, page 17-33

ICA Acceleration Report

The ICA Acceleration report displays the ICA acceleration statistics. The following charts and tables are included:

- ICA: Effective WAN Capacity, page 17-24
- ICA: Connection Details, page 17-24
- ICA: Client Versions, page 17-24
- ICA: Unaccelerated Reasons, page 17-25
- ICA Acceleration Statistics Table, page 17-31

Note The ICA charts in WAAS version 5.0 and later are different from those used in version 4.5. If you are viewing the data from a version 4.5 WAAS device, the charts appear empty due to the different data that the device is collecting. The ICA data for version 4.5 WAAS devices is available in the system level TCP Summary Report.

Summary Report

The Summary Report is a predefined report that can be used to monitor the system operation. It is available at the system level. This report displays the following charts and tables by default:

- Traffic Summary, page 17-16
- Effective WAN Capacity, page 17-15
- Traffic Summary Over Time, page 17-16
- Traffic Volume and Reduction, page 17-16
- Compression Summary, page 17-15
- Compression Summary Over Time, page 17-15
- HTTP: Estimated Time Savings, page 17-17
- HTTP: Effective WAN Capacity, page 17-17
- MAPI: Effective WAN Capacity, page 17-21
- SSL: Effective WAN Capacity, page 17-20
- MAPI: Average Response Time Saved, page 17-21
- Network Application Traffic Details Table, page 17-30

The Summary Report can be customized to display the charts that you require. Use the Customize taskbar icon to select the charts that you want to be displayed on this report. Only 12 charts can be displayed in the report.

Topology Report

The Topology report at the system level displays a topology map that shows a graphical representation of all the connections between the WAAS devices.

The topology map uses blue squares to show connections between devices. Use the legend to the right of the grid to associate the device name with the number that appears at the top of the grid. Use the drop-down lists at the top of the window to perform the following tasks:

- Display connections between your various locations instead of between devices.
- Sort the grid by the number of connections instead of by device name.

Click the **View** icon next to the WAE to view a list of peer devices for a specific WAE. The Peer List window appears, which is the same as the device level Topology report.

At the device level, the Topology report lists all the peer devices connected to a specific WAE so that you can see the relationship between devices in your WAAS network. The Peer List window displays information about each peer device involved in optimized connections with this WAE. To go to the system level Topology report, click the **Topology** icon in the taskbar.

If a peer device is not registered with the WAAS Central Manager, the message “Unknown, this peer is not being managed by CM” is shown for the name and “Unknown” is displayed for the IP address.

Note The WAAS Central Manager device does not have any peers because it does not participate with any WAEs to optimize traffic. For this reason, the topology feature is not available on the WAAS Central Manager device.

Connection Trend Report

The Connection Trend Report displays the connection trends of applications on a device. The following charts are included:

- Optimized Connections Over Time, page 17-26
- Optimized vs Pass-Through Connections, page 17-26

Connections Statistics Report

The Connections Statistics report displays a Connections Summary table for the device. The table displays all the TCP connections handled by the device and corresponds to the **show statistics connection** EXEC mode command in WAE and the **show waas connection brief** command in WAAS Express.

The report displays the following information about each connection:

- Source IP address and port.
- Destination IP address and port.
- Peer ID—Hostname of the peer device.

- Applied Policy/Bypass Reason—Displays icons representing the applied optimization policies, including TFO, DRE, LZ, and an application accelerator, respectively (hover your mouse over the icon to see its meaning). If the connection was not optimized, the bypass reason is shown.
- Connection Start Time—Date and time when the connection was started.
- Open Duration—Number of hours, minutes, and seconds that the connection has been open.
- Total number of original bytes.
- Total number of optimized bytes.
- Percentage of compression.
- Classifier name—If no classifier exists for the connection, this column contains a Create New button. Click the button to display a Classifier Settings form below the table where you can create a classifier that matches the source and destination IP addresses and ports of the connection. Enter a name in the Classifier Name field, check the **Match All** check box to match all traffic, or make selections from the Source IP, Source Port, Destination IP, and Destination Port drop-down lists. Then, click the **Create Classifier** button to create the classifier.

Note If the WAE is inheriting policies from a device group, the Create New button is not shown, to prevent a user from unknowingly overriding device group policies. To create a classifier, you must first override the device group policy page and then return to the Connection Statistics report.

The data in the Connections Summary Table is retrieved from the device one time when you view the window for the first time.

Click the **Refresh** button at the bottom of the window to refresh the data in the Connections Summary Table.

From the Connections Summary Table for Device window, you may perform the following tasks:

- Apply filter settings to display particular connections based on criteria that you choose.
- View connection details.
- Click the **Reset Filter** button to reset the filter options and refresh the table.

Click the **Details** icon next to the connection entry in the summary table to view connection details. The Connection Details window appears. This window contains connection addresses, port information, policy information, and traffic statistics. The Connection Details window also displays graphs that plot real-time traffic statistics and are refreshed every two seconds.

Note If the value for Percentage Compression is negative, the Percentage Compression and Effective Capacity values do not appear.

CIFS Acceleration Report

The CIFS Acceleration report displays the CIFS acceleration statistics for a WAAS device. The following charts are included on two tabs:

- CIFS: Connection Statistics—Displays the number of CIFS accelerated sessions.
- CIFS: File Optimization—Displays the number of open CIFS files.
- CIFS: Request Optimization—Displays the percentage of requests that are served locally from the CIFS cache.

- CIFS: Cache Utilization—Displays the utilization percentage of the CIFS cache.
- CIFS: Cached Objects—Displays the number of objects in the CIFS cache.
- CIFS: Client Average Throughput—Displays the average throughput (in KB/second) between the WAAS device and its clients.

Note When you use the Print icon in the taskbar to print the CIFS Acceleration report to a file, all the CIFS charts will display the time in WAE local time (the CE Local Time setting), regardless of the chart time zone settings that you have configured.

CIFS Acceleration Report for WAAS Express

The CIFS Acceleration report for WAAS Express displays the CIFS acceleration statistics for a WAAS Express device. The following charts and tables are included:

- CIFS: Client Average Throughput, page 17-25
- CIFS: Connection Details, page 17-25
- CIFS: Effective WAN Capacity, page 17-25
- CIFS: Request Optimization, page 17-25
- CIFS Acceleration Statistics Table for WAAS Express, page 17-34

Resource Utilization Report

The Resource Utilization report displays the following charts:

- CPU Utilization
- Disk Utilization

Disks Report

The Disks Report displays physical and logical disk information.

The report window displays the following information about each disk:

- Physical disk information, including the disk name, serial number, and disk size.
- Present status. The Present field will show either Yes if the disk is present or Not Applicable if the disk is administratively shut down.
- Operational status (NORMAL, REBUILD, BAD, UNKNOWN, or Online).
- Administrative status (ENABLED or DISABLED). When the Administrative Status field shows DISABLED, the Present field will show Not Applicable.
- Current and future disk encryption status.
- Current and future extended object cache status.
- RAID level. For RAID-5 devices, the Disk Information window includes the RAID device name, RAID status, and RAID device size.
- Error information, if any errors are detected.

From this window, you may save all disk information details to an Excel spreadsheet by clicking the **Export Table** icon in the taskbar.

AppNav Report

The AppNav report displays AppNav flow distribution information. This report is available at the AppNav Cluster level, where it shows statistics for the whole AppNav Cluster, and at the device level for AppNav Controllers (ANCs), where it shows statistics for a single ANC.

The following charts and tables are included:

- Total AppNav Traffic, page 17-26
- AppNav Policies, page 17-26
- Top 10 AppNav Policies, page 17-27
- Top 10 WAAS Node Group Distribution, page 17-27
- WAAS Node Group Distribution, page 17-27
- Pass-Through Reasons, page 17-27
- Top 10 Pass-Through Reasons, page 17-27

At the AppNav Cluster level, the following additional controls appear in the taskbar:

- The Scope pull-down list allows you to choose to display data for the whole cluster or for an individual ANC.
- The AppNav Policy Rule pull-down list allows you to choose the AppNav policy for which data is displayed.

Note At the AppNav Cluster level, the charts may not show data if the configuration on all ANCs in the cluster does not match. To resolve this situation, choose **AppNav Clusters >** *cluster-name* from the Central Manager menu and click the taskbar icon named Force Settings on all Devices in a Group. After about 15 minutes, the AppNav charts should display data.

Managing Reports

The WAAS Central Manager allows you to edit any of the predefined reports and to create custom reports. Additionally, you can schedule reports to be generated periodically such as hourly, daily, weekly, or monthly. When a scheduled report is generated, a link to the report is e-mailed to notify the recipients.

This section contains the following topics:

- Creating Custom Reports, page 17-44
- Viewing and Editing Reports, page 17-45
- Scheduling Reports, page 17-46
- Managing Scheduled Reports, page 17-47

Creating Custom Reports

A report consists of up to eight charts and tables. The system and device dashboard displays are examples of predefined reports, along with the other reports available in the Monitor menu.

Reports can be created only at the system level, not at the device level.

To create a custom report, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Monitor > Reports > Reports Central**.

Step 2 Click the **Create** taskbar icon. The Create Report pane appears, as shown in Figure 17-8.

Figure 17-8 Create Report Pane

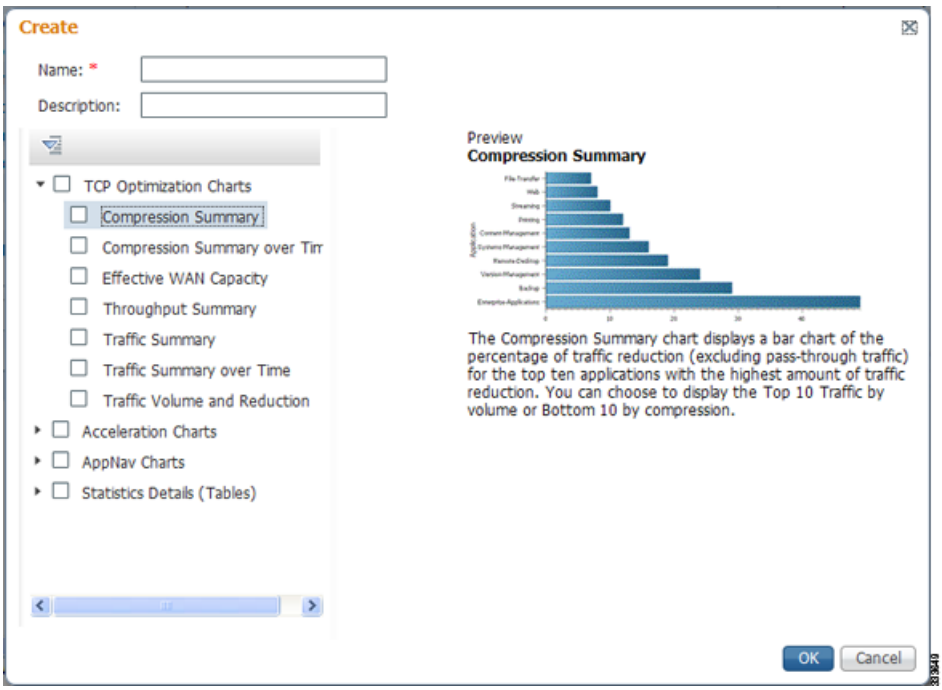

Step 3 In the Name field, enter a name for the report using up to 64 characters. Only the following characters are allowed to be entered: numbers, letters, spaces, periods, hyphens, and underscores.

Step 4 (Optional) In the Description field, enter a description of the report.

Step 5 In the list at the left side of the pane, check the box next to each chart and table that you want to display in the report. See the "Chart and Table Descriptions" section on page 17-14 for a description of the charts.

Expand any of the categories by clicking on the small triangle next to the category name. See a preview and description of any chart by clicking on the chart name. Tables are listed in the last category, Statistics Details.

Step 6 Click **OK**.

Step 7 (Optional) Customize any of the chart settings as follows:

a. Display the report by clicking the report name in the Report Templates table. You may have to scroll the table.

b. You can customize report settings such as the time frame and the time zone as described in the "Customizing a Dashboard or Report" section on page 17-10.

c. Click the **Edit** icon in the upper left of a chart to customize the chart settings. For more information, see the "Configuring Chart Settings" section on page 17-14.

d. Click **OK**.

Repeat the steps for each chart you want to customize.

Another way you can create a report is to copy a similar existing report and modify it into a new report. To copy a report, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Monitor > Reports > Reports Central**.

Step 2 Check the box next to the report that you want to copy.

Step 3 Click the **Copy** taskbar icon. The copy report window appears.

Step 4 In the Name field, enter a name for the report.

Step 5 (Optional) In the Description field, enter a description of the report.

Step 6 Click **OK**.

The report is added to the Reports table.

Viewing and Editing Reports

To view or edit a report, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Monitor > Reports > Reports Central**.

Step 2 Click the name of the report that you want to view or edit.

If you do not see the report that you are looking for, you may need to scroll the Reports table. You can filter the list by choosing **Quick Filter** from the Show drop-down list and entering filter criteria.

Step 3 If you want to change any of the charts or tables in the report, use the standard chart editing methods as described in the "Customizing a Dashboard or Report" section on page 17-10.

Step 4 Click **Save** or **Save As** to save the report.

To delete a report from the Reports table, check the check box next to the report and click the **Delete** taskbar icon.

Admin users can view, edit, and delete reports created by all users and can view and edit predefined reports. Non-admin users can view, edit, and delete only reports created by themselves and can view and edit predefined reports.

Scheduling Reports

You can schedule reports to be generated once or periodically such as daily, weekly, or monthly. When a scheduled report is generated, a copy of the report can be e-mailed.

To schedule a report, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Monitor > Reports > Reports Central**.

Step 2 Check the box next to the report that you want to schedule.

If you do not see the report that you are looking for, you may need to scroll the Reports table.

Step 3 Click the **Schedule** icon in the taskbar. The scheduling window appears, as shown in Figure 17-9.

Figure 17-9 Scheduling a Report

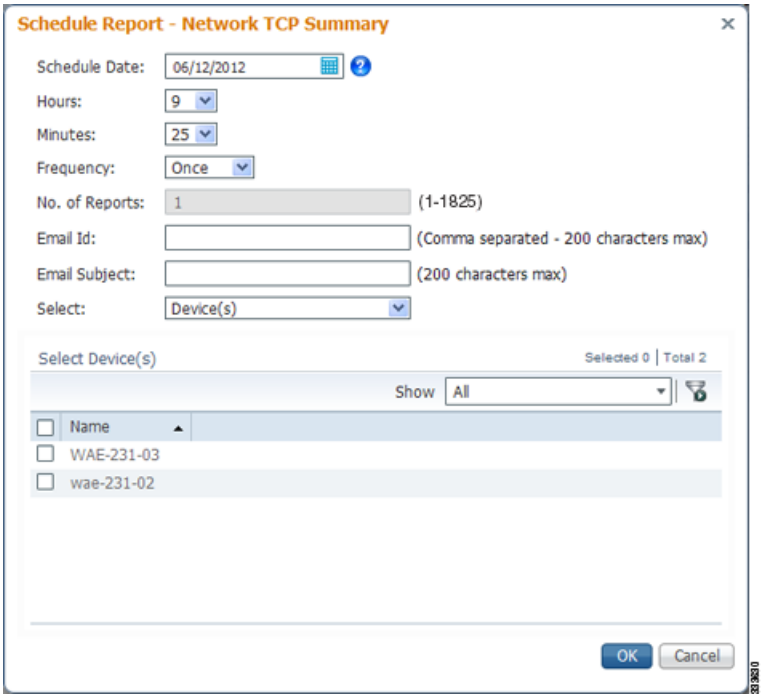

Step 4 In the Date field, enter the schedule date in the format DD/MM/YYYY or click the calendar icon to display a calendar popup window from which to choose the date.

Step 5 In the Hours drop-down list, choose the hours. The time represents the local time at the WAAS Central Manager.

Step 6 In the Minutes drop-down list, choose the minutes. The time represents the local time at the WAAS Central Manager.

Step 7 In the Frequency drop-down list, choose Once, Hourly, Daily, Weekly, or Monthly for the report frequency.

Step 8 In the No. of Reports field, enter the number of times that a reoccurring report is to be generated. After being generated the specified number of times, the report is no longer generated.

Step 9 In the Email Id field, enter the e-mail addresses of the report recipients, separated by commas.

Step 10 In the Email Subject field, enter the subject of the e-mail message.

Step 11 In the Select drop-down list, choose **Device(s)**, **DeviceGroup**, **Cluster**, or **Location** to display a list of the chosen entities.

Step 12 In the Select entity area, choose the devices that are to be included in the statistics for the report. Place a check in the box next to each device, device group, cluster, or location that you want to include.

To find (highlight) an entity in a long list, choose **Quick Filter** from the Show drop-down list and enter the entity name (or partial name) in the field above the list. The search is case sensitive.

Step 13 Click **OK**.

Step 14 Configure the e-mail server settings for e-mail notification when reports are generated. For more information, see the "Configuring the E-mail Notification Server" section on page 10-24.

Note In a WAAS network where there are 1000 or more WAEs, a scheduled report might take up to 4 minutes to generate. And if you schedule more than one report at the same time, the reports will be generated with a delay of up to 20 minutes, depending on the number of reports and devices.

Managing Scheduled Reports

To view or delete a scheduled report, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Monitor > Reports > Reports Central**.

The lower part of the Reports window lists the completed and pending scheduled reports, depending on which tab you choose. You can use the Show filter above the table to filter the reports that are displayed.

Step 2 (Optional) If you want to view a completed report instance in the Completed Reports tab, click the **Completed** link in the Status column.

Note For each completed instance of a scheduled report, the Frequency column shows Once and the Completed Time shows the date and time that the report was generated.

Step 3 (Optional) If you want to view a list of pending reports, click the **Pending Reports** tab.

Step 4 (Optional) If you want to delete a report in either the Completed Reports or Pending Reports tabs, check the box next to one or more report instances that you want to delete and click the **Delete** taskbar icon.

WAAS stores the 10 most recently completed or failed report instances for each custom report. This number is configurable by the System.monitoring.maxReports system property. For details on changing this property, see the "Modifying the Default System Configuration Properties" section on page 10-17.

Admin users can view reports scheduled by all users and the name of the report creator. Non-admin users can view only reports scheduled by themselves.

Any changes to predefined report settings are stored separately for individual users. That is, if one user changes a predefined scheduled report, only that user sees the changes, and other users (including admin users) continue to see the report with default settings.

Any reports scheduled by an external user are deleted if the maximum limit of days without a login passes and the user is deleted. For more information, see the cdm.remoteuser.deletionDaysLimit system configuration property in Table 10-4 on page 10-18.

Configuring Flow Monitoring

Flow monitoring applications collect traffic data that is used for application trend studies, network planning, and vendor-deployment impact studies. This section describes how to configure the flow monitoring feature on the WAE and includes the following topics:

- Alarms for Flow Monitoring, page 17-49
- Example Using NetQoS for Flow Monitoring, page 17-50

The NetQoS monitoring application can interoperate with the WAAS software to provide flow monitoring. To integrate this application with the WAAS software, you configure the NetQoS FlowAgent module on the WAE devices. The NetQoS FlowAgent module on the WAE collects important metrics of packet flows, which are then sent across the network to the NetQoS SuperAgent. This monitoring agent analyzes the data and generates reports. For this feature to work, additional configuration is required on the NetQoS FlowAgent. (See the "Example Using NetQoS for Flow Monitoring" section on page 17-50.)

The monitoring agent is composed of two modules: the console (or host) and the collector. The WAE initiates two types of connections to these two monitoring agent modules: a temporary connection to the console and a persistent connection to the collector. You configure the console IP address on the WAE by entering the **flow monitor tcpstat-v1 host** configuration mode command in either the WAE CLI or through the Central Manager GUI. This temporary connection is referred to as the control connection. The control connection uses TCP port 7878. Its purpose is to obtain the IP address and port number of the collector to which the WAE is assigned. The WAE also pulls the configuration information regarding which servers are to be monitored over the control connection. Once the WAE obtains the IP address and port number of the collector, the WAE opens a persistent connection to the collector. Collected summary data for the servers that are being monitored is sent over this persistent connection.

You may place the console (or host) module and the collector module on a single device or on separate devices. These connections are independent of one another. A failure of one connection does not cause the failure of the other connection and vice versa.

The state of these connections and various operation statistics display when you use the **show statistics flow monitor tcpstat-v1** EXEC mode command. Connection errors and data transfer errors trigger alarms on the WAE and in the Central Manager GUI. (See the "Alarms for Flow Monitoring" section on page 17-49.) To display debug information, use the **debug flow monitor tcpstat-v1** EXEC mode command.

To configure flow monitoring on your WAEs using the Central Manager GUI, follow these steps:

Step 1 Create a new device group for configuring flow monitoring on multiple devices. Choose **Device Groups >** *device-group-name* **> Create New Device Group** to create a device group.

a. When you create the device group, check the **Automatically assign all newly activated devices to this group** check box to enable this option.

b. Add your existing WAE devices to this new device group.

Step 2 From the Device Group listing window, click the **Edit** icon next to the name of the flow monitoring configuration device group that you want to configure.

Step 3 Choose **Configure > Monitoring > Flow Monitor**. The Flow Monitor Settings for Device Group window appears.

Step 4 Check the **Enable** check box.

Step 5 In the tcpstat-v1 Host field, enter the IP address of the monitoring agent console.

This configuration allows the WAE to establish a temporary connection (a control connection) to the console for the purpose of obtaining the IP address of the collector device. You must configure the collector IP address information from the console device. (See the configuration documentation for the NetQoS flow monitoring application software.)

Step 6 Click **Submit** to apply the settings to the devices in this device group.

To configure flow monitoring on the WAE using the CLI, follow these steps:

Step 1 Register the WAE with the IP address of the monitoring agent console.

```
WAE(config)# flow monitor tcpstat-v1 host 10.1.2.3
```

This configuration allows the WAE to establish a temporary connection (a control connection) to the console (or host) for the purpose of obtaining the IP address of the collector device. You must configure the collector IP address information from the console device. (See the configuration documentation for the NetQoS flow monitoring application software.)

Step 2 Enable flow monitoring on the WAE appliance.

```
WAE(config)# flow monitor tcpstat-v1 enable
```

Step 3 Check the configuration by using the **show running-config** EXEC command.

Alarms for Flow Monitoring

Table 17-15 describes the four different alarms that may be raised when errors occur with flow monitoring.

Table 17-15 Alarms for Flow Monitoring

Name	Severity	Description
CONTROL_CONN	Major	Indicates a problem with the control connection.
COLLECTOR_CONN	Major	Indicates a problem with the collector connection.
SUMMARY_COLLECTION	Minor	Indicates a problem with the collection of packet summary information. Summary packets may be dropped because the buffer queue limit has been reached or because of a TFO error, such as not being able to allocate memory. Summary packet collection may also be dependent on available WAN bandwidth.
DATA_UPDATE	Minor	Indicates a problem with the ability of the WAE to send updates the collector agent.

Example Using NetQoS for Flow Monitoring

NetQoS integrates with the WAAS software by running the NetQoS FlowAgent on WAE devices. FlowAgent is a software module developed by NetQoS that resides on the WAE appliance. The FlowAgent collects metrics about the packet flows, which are then sent across the network to a NetQoS SuperAgent. The SuperAgent measures the round-trip times, server response times, and data transfer times, and then analyzes the data and generates reports.

Note When you use flow monitoring with the NetQoS SuperAgent, the flow monitor on the WAE captures optimized traffic only.

To configure flow monitoring with NetQoS, follow these steps:

Step 1 From the WAE CLI or Central Manager GUI, enter the SuperAgent Master Console IP address in the tcpstat-v1 Host field on your WAE appliances.

If you are configuring multiple appliances through a device group, wait for the configuration to propagate to all the appliances in the device list.

Step 2 From the NetQoS SuperAgent console, assign a WAE to a SuperAgent Aggregator (known as the collector in WAAS terminology) and configure the NetQoS Networks, Servers, and Applications entities.

Note For information about using the NetQoS SuperAgent Master Console and configuring NetQoS SuperAgent entities, go to the following website: http://support.ca.com

Configuring and Viewing Logs

This section contains the following topics:

Configuring System Logging

Use the WAAS system logging feature to set specific parameters for the system log file (syslog). This file contains authentication entries, privilege level settings, and administrative details. The system log file is located on the system file system (sysfs) partition as /local1/syslog.txt.

To enable system logging, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Monitoring > Log Settings > System Log**. The System Log Settings window appears. (See Figure 17-10.)

Figure 17-10 System Log Settings Window

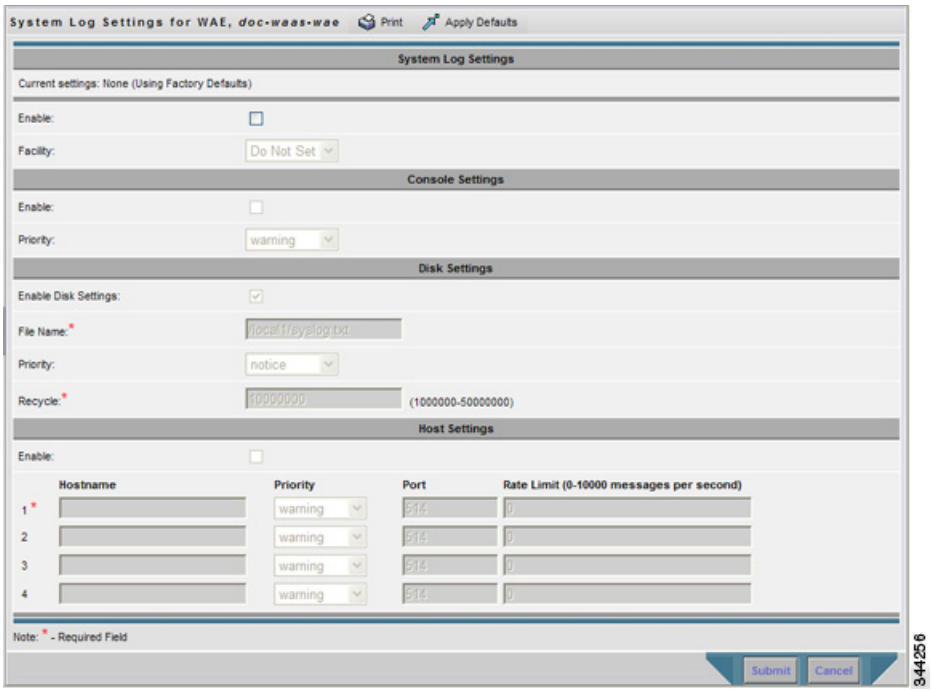

Step 3 Under the System Log Settings section, check the **Enable** check box to enable system logging. By default, this option is disabled.

Step 4 From the Facility drop-down list, choose the appropriate facility.

Step 5 Enable system log files to be sent to the console:

a. In the Console Settings section, check the **Enable** check box.

b. From the Priority drop-down list, choose the severity level of the message that should be sent to the specified remote syslog host. The default priority-code is “warning” (level 4). Each syslog host is capable of receiving a different level of event messages. (See Table 17-16 on page 17-52 for a list of priority levels.)

Step 6 Enable syslog files to be sent to a disk:

a. In the Disk Settings section, check the **Enable Disk Settings** check box.

b. In the File Name field, enter a path and a filename where the syslog files will be stored on a disk.

c. From the Priority drop-down list, choose the severity level of the message that should be sent to the specified remote syslog host. The default priority code is “warning” (level 4). Each syslog host is capable of receiving a different level of event messages. (See Table 17-16 on page 17-52 for a list of priority levels.)

d. In the Recycle field, specify the size of the syslog file (in bytes) that can be recycled when it is stored on a disk. The default value of the file size is 10000000.

   Whenever the current log file size surpasses the recycle size, the log file is rotated. (The default recycle size for the log file is 10,000,000 bytes.) The log file cycles through at most five rotations, and each rotation is saved as *log_file_name.[1-5]* under the same directory as the original log.

   The rotated log file is configured in the File Name field (or by using the **logging disk filename** command).

Step 7 Enable syslog files to be sent to a host:

a. In the Host Settings section, check the **Enable** check box. You can configure up to four hosts to which syslog messages can be sent. For more information, see the "Multiple Hosts for System Logging" section on page 17-53.

b. In the Hostname field, enter a hostname or IP address of the remote syslog host. Specify up to three more remote syslog hosts in the Hostname fields 2 through 4. You must specify at least one hostname if you have enabled system logging to a host.

c. From the Priority drop-down list, choose the severity level of the message that should be sent to the specified remote syslog host. The default priority code is "warning" (level 4). Each syslog host is capable of receiving a different level of event messages. (See Table 17-16 for a list of priority levels.)

d. In the Port field, specify the destination port on the remote host to which the WAAS device should send the message. The default port number is 514.

e. In the Rate Limit field, specify the number of messages per second that are allowed to be sent to the remote syslog host. To limit bandwidth and other resource consumption, messages to the remote syslog host can be rate limited. If this limit is exceeded, the specified remote syslog host drops the messages. There is no default rate limit, and by default all syslog messages are sent to all of the configured syslog hosts.

Step 8 Click **Submit**.

To configure system logging from the CLI, you can use the **logging** global configuration command.

This section contains the following topics:

Priority Levels

Table 17-16 lists the different priority levels of detail to send to the recipient of the syslog messages for a corresponding event.

Table 17-16 System Logging Priority Levels and Descriptions

Priority Code	Condition	Description
0	Emergency	System is unusable.
1	Alert	Immediate action needed.
2	Critical	Critical condition.
3	Error	Error conditions.

Table 17-16 System Logging Priority Levels and Descriptions (continued)

Priority Code	Condition	Description
4	Warning	Warning conditions.
5	Notice	Normal but significant conditions.
6	Information	Informational messages.
7	Debug	Debugging messages.

Multiple Hosts for System Logging

Each syslog host can receive different priority levels of syslog messages. You can configure different syslog hosts with a different syslog message priority code to enable the WAAS device to send varying levels of syslog messages to the four external syslog hosts. For example, a WAAS device can be configured to send messages that have a priority code of "error" (level 3) to the remote syslog host that has an IP address of 10.10.10.1 and messages that have a priority code of "warning" (level 4) to the remote syslog host that has an IP address of 10.10.10.2.

If you want to achieve syslog host redundancy or failover to a different syslog host, you must configure multiple syslog hosts on the WAAS device and assign the same priority code to each configured syslog host (for example, assigning a priority code of "critical" (level 2) to syslog host 1, syslog host 2, and syslog host 3).

In addition to configuring up to four logging hosts, you can also configure the following for multiple syslog hosts:

- A port number different from the default port number, 514, on the WAAS device to send syslog messages to a logging host.
- A rate limit for the syslog messages, which limits the rate at which messages are sent to the remote syslog server (messages per second) to control the amount of bandwidth used by syslog messages.

Configuring Transaction Logging

This section contains the following topics:

Enabling Transaction Logging

To enable transaction logging for TFO flows and video streams, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Monitoring > Log Settings > Transaction Log** for TFO transaction logging or **Configure > Monitoring > Log Settings > Video Acceleration Transaction Log** for video transaction logging. The Transaction Log Settings window appears. (See Figure 17-11.) The Video Transaction Log Settings window looks the same, but does not include the General Settings area at the top.

Figure 17-11 Transaction Log Settings Window

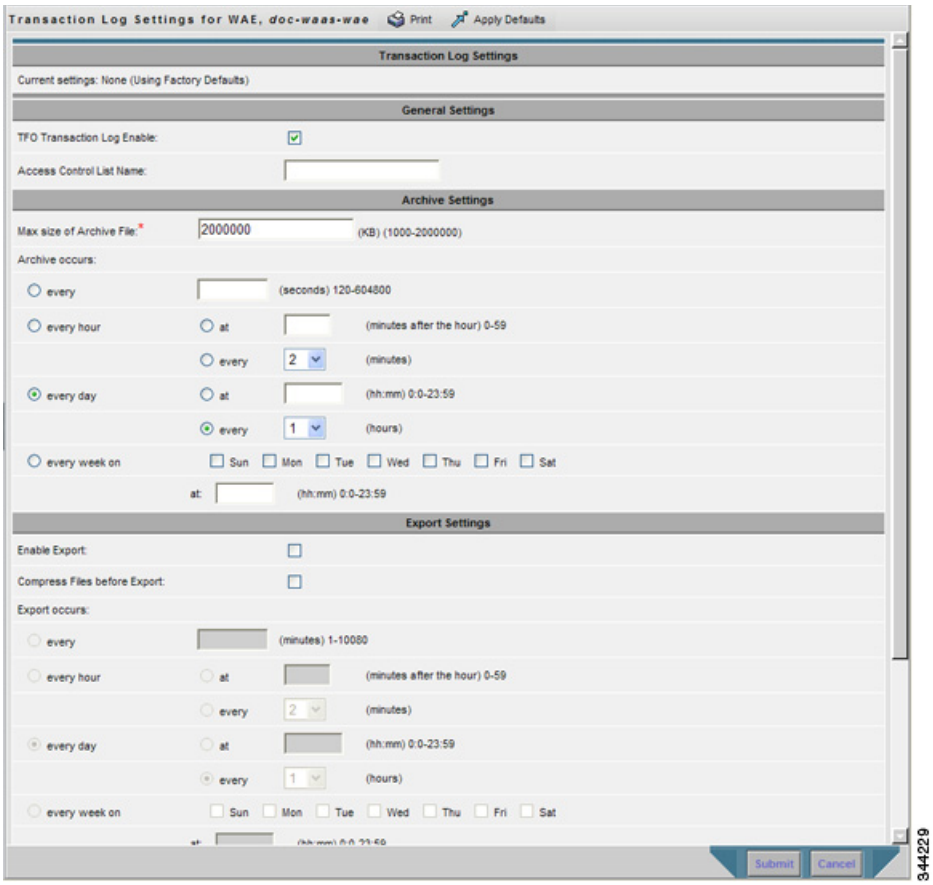

Step 3 Under the General Settings heading, check the **TFO Transaction Log Enable** check box to enable transaction logging. This check box does not appear for video transaction logging.

The fields on the window become active.

Step 4 In the Access Control List Name field, optionally enter the name of an access control list that you want to use to limit transaction logging. If you specify an access control list, only transactions from hosts that are defined in the access list are logged. This field does not appear for video transaction logging.

Use the **ip access-list** global configuration command to define an access list.

Step 5 Under the Archive Settings heading, specify values for the following fields:

- Max Size of Archive File—Maximum size (in kilobytes) of the archive file to be maintained on the local disk. This value is the maximum size of the archived file to be maintained on the local disk. The range is 1000 to 2000000. The default is 2000000.
- Archive Occurs Every (interval)—Interval at which the working log data is cleared and moved into the archive log.

Step 6 Configure the fields in the Export Settings section to export the transaction log file to an FTP server.

Table 17-17 describes the fields in the Export Settings section.

Table 17-17 Export Settings

Field	Function
Enable Export	Enables transaction logging to be exported to an FTP server.
Compress Files before Export	Enables compression of archived log files into gzip format before exporting them to external FTP servers.
Export occurs every (interval)	Interval at which the working log should be cleared by moving data to the FTP server.
Export Server	The FTP export feature can support up to four servers. Each server must be configured with a username, password, and directory that are valid for that server. • Export Server—The IP address or hostname of the FTP server. • Name—The user ID of the account used to access the FTP server. • Password/Confirm Password—The password of the FTP user account specified in the Name field. You must enter this password in both the Password and Confirm Password fields. Do not use the following characters: space, backward single quote (`), double quote ("), pipe (|), or question mark (?). • Directory—The name of a working directory that will contain the transaction logs on the FTP server. The user specified in the Name field must have write permission to this directory. • SFTP—If the specified FTP server is a secure FTP server, check the **SFTP** check box.

Step 7 Click **Submit**.

A "Click Submit to Save" message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking **Reset**. The Reset button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

If you try to leave this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box only appears if you are using the Internet Explorer browser.

To enable and configure transaction logging from the CLI, you can use the **transaction-logs** global configuration command.

Transaction Logs

TFO transaction logs are kept on the local disk in the directory /local1/logs/tfo. Video (Windows media) logs are kept in the directory /local1/logs/wmt/wms-90.

When you enable transaction logging, you can specify the interval at which the working log should be archived by moving the data to an archive log. The archive log files are located on the local disk in the directory /local1/logs/.

Because multiple archive files are saved, the filename includes the time stamp when the file was archived. Because the files can be exported to an FTP/SFTP server, the filename also contains the IP address of this WAAS device.

The archive filenames for TFO transactions use this format:

tfo_IPADDRESS_YYYYMMDD_HHMMSS.txt.

The archive filenames for Windows media transactions use this format:

wms_90_IPADDRESS_YYYYMMDD_HHMMSS.txt.

The transaction log format is documented in Appendix B, "Transaction Log Format."

Viewing the System Message Log

Using the system message log feature of the WAAS Central Manager GUI, you can view information about events that have occurred in your WAAS network. The WAAS Central Manager logs messages from registered devices with a severity level of "warning" or higher.

To view logged information for your WAAS network, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > Logs > System Messages**. The System Message Log window appears. (See Figure 17-12.)

Figure 17-12 System Message Log

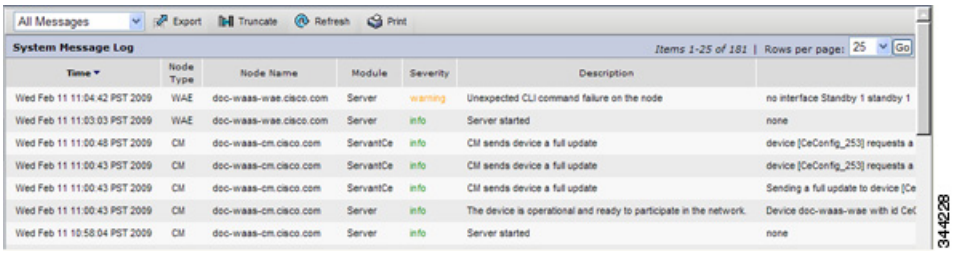

Step 2 From the System Message Log drop-down list, choose one of the following types of messages to display:

- **All**
- **CLI**
- **Critical**
- **Database**

Step 3 (Optional) Click a column heading by node type, node name, module, or message text to sort the messages. By default, messages are listed chronologically.

Note If no name is available for a node, the name displayed is "Unavailable." This situation might occur if the node has been deleted or has been reregistered with WAAS software.

Step 4 (Optional) Truncate the message log so that not as many messages appear in the table, by completing the following steps:

a. Click the **Truncate** icon in the taskbar. The Truncate System Message Log window appears.

b. Choose one of the following options:

- **Size Truncation**—Limits the messages in the log to the number you specify. The log uses a first in, first out process to remove old messages once the log reaches the specified number.
- **Date Truncation**—Limits the messages in the log to the number of days you specify.
- **Message Truncation**—Removes messages from the log that match the specified pattern.

c. Click **Submit** when you have finished specifying the truncation parameters.

Step 5 If you have many event messages, you may need to view multiple pages to view the activity in which you are interested. Click the forward (>>) and back (<<) buttons to move between pages. Alternatively, click the link for a specific page number to jump to that page.

Viewing the Audit Trail Log

The WAAS Central Manager logs user activity in the system. The only activities that are logged are those activities that change the WAAS network. This feature provides accountability for users' actions by describing the time and action of the task. Logged activities include the following:

- Creation of WAAS network entities
- Modification and deletion of WAAS network entities
- System configurations
- Clearing the audit log

To view audit trail logs, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Admin > Logs > Audit Trail Logs**.

The Audit Log window appears. (See Figure 17-13.) All logged transactions in the WAAS Central Manager are listed by date and time, user, actual transaction that was logged, and the IP address of the machine that was used.

Figure 17-13 Audit Log Window

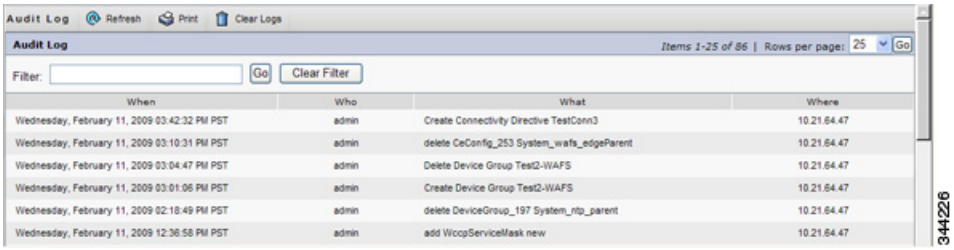

Step 2 Choose a number from the Rows drop-down list to determine the number of rows that you want to display.

Viewing the Device Log

To view information about events that have occurred on a specific device in your WAAS network, you can use the system message log feature available in the WAAS Central Manager GUI.

To view events that have occurred on your entire WAAS network, see the "Viewing the System Message Log" section on page 17-56.

To view the logged information for a WAAS device, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Admin > History > Logs**. The System Message Log for Device window appears.

Step 3 Choose the type of messages to be displayed from the System Message Log drop-down list.

Choose one of the following types of messages in the system log:

- **All** (default)
- **CLI**
- **Critical**
- **Database**

Step 4 Click a column heading to arrange the messages chronologically by node type, node name, or module. By default, messages are displayed chronologically.

If no name is available for a node because the node has been deleted or reregistered with the WAAS software, the message displayed is "Unavailable."

Step 5 If you have many event messages, you may need to use the forward (>>) and back (<<) buttons to move between pages. Alternatively, click the link for a specific page number to move to that particular page.

Troubleshooting Tools

This section contains the following topics:

For additional advanced WAAS troubleshooting information, see the *Cisco WAAS Troubleshooting Guide for Release 4.1.3 and Later* on Cisco DocWiki.

Enabling the Kernel Debugger

The WAAS Central Manager GUI allows you to enable or disable access to the kernel debugger (kdb). Once enabled, the kernel debugger is automatically activated when kernel problems occur.

To enable the kernel debugger, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Monitor > Tools > Kernel Debugger**. The Kernel Debugger window appears.

Step 3 Check the **Enable** check box to enable the kernel debugger, and click **Submit**. By default, this option is disabled.

Using Diagnostic Tests

WAAS includes diagnostic testing tools as described in the following sections:

- Diagnostic Testing Using the GUI, page 17-59
- Diagnostic Testing Using the CLI, page 17-60

Diagnostic Testing Using the GUI

The WAAS Central Manager includes a troubleshooting and diagnostic reporting facility.

To perform diagnostic tests, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Monitor > Tools > Diagnostics Tests**. The Diagnostic Tool window appears.

Step 3 Check the check box next to each diagnostic test that you want to run, or check the top check box to run all tests. The following tests are available:

- Device Operation—Check the device status and the presence of coredump files or alarms of major or critical severity.
- Basic Configuration—Check the device basic network configuration.
- Basic Connectivity—Check the device connectivity to configured external devices (DNS, authentication, NTP servers, and so forth).
- Physical Interface—Check the configuration and operation of device physical interfaces.

  Note A Virtual Interface test is available for vWAAS devices.

- Configuration Security—Check the running configuration for potentially malicious (XSS) entries.
- Traffic Optimization—Check the TFO configuration and operation.
- WCCP configuration and operation—Check the configuration and operation of WCCP traffic interception.
- Inline configuration and operation—Check the configuration and operation of inline group interfaces.

  Note The inline configuration and operation test is not available for vWAAS devices.

Step 4 Click **Run**.

Step 5 View the test results in the lower part of the window. You may have to scroll the window to see all results.

For tests that fail, error messages describe the problem and provide recommended solutions.

You can run the same diagnostic tests again and refresh the results by clicking the **Refresh** icon in the taskbar.

To print the results, click the **Print** icon in the taskbar.

Diagnostic Testing Using the CLI

You can use the **test** EXEC command to perform diagnostic and connectivity tests.

You can use network-level tools to intercept and analyze packets as they pass through your network. Two of these tools are TCPdump and Tethereal, which you can access from the CLI by using the **tcpdump** and **tethereal** EXEC commands.

The WAAS device also supports multiple debugging modes, reached with the **debug** EXEC command. These modes allow you to troubleshoot problems from configuration errors to print spooler problems. We recommend that you use the **debug** command only at the direction of Cisco TAC.

The output associated with the **debug** command is written to either the syslog file in /local1/syslog.txt or the debug log associated with the module in the file /local1/errorlog/*module_name*-errorlog.current.

The output associated with the **debug accelerator** *name module* command for an application accelerator is written to the file *name*ao-errorlog.current, where *name* is the accelerator name. The accelerator information manager debug output is written to the file aoim-errorlog.current.

The debug log file associated with a module will be rotated to a backup file when the current file reaches its maximum size. The backup files are named as follows: *name*-errorlog.#, where # is the backup file number.

For any **debug** command, system logging must be enabled. The command to enable logging is the **logging disk enable** global configuration command, which is enabled by default.

If a **debug** command module uses the syslog for debug output, the **logging disk priority debug** global configuration command must be configured (the default is **logging disk priority notice**).

If a **debug** command module uses the debug log for output, the output can be filtered based on a priority level configuration for the four different levels of debug log output, as follows:

- For filtering on critical debug messages only, use the following global configuration command: **logging disk priority critical**.
- For filtering on critical and error level debug messages, use the following global configuration command: **logging disk priority error**.
- For filtering on critical, error, and trace debug level debug messages, use the following global configuration command: **logging disk priority debug**.
- For seeing all debug log messages, which include critical, error, trace and detail messages, use the following global configuration command: **logging disk priority detail**.

Regardless of the priority level configuration, any syslog messages at the LOG_ERROR or higher priority will be automatically written to the debug log associated with a module.

For more details on these CLI commands, see the *Cisco Wide Area Application Services Command Reference*.

Using the show and clear Commands from the WAAS Central Manager GUI

To use the WAAS Central Manager GUI **show** and **clear** command tool, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*.

Step 2 Choose **Monitor > CLI Commands > Show Commands** or **Clear Commands**.

Step 3 From the drop-down list, choose a **show** or **clear** command.

Step 4 Enter arguments for the command, if any.

Step 5 Click **Submit** to display the command output.

A window appears, displaying the command output for that device.

The **show** and **clear** CLI commands that are available differ depending on the type of device that is selected.

You can also use the **show** EXEC commands from the CLI. For more information, see the *Cisco Wide Area Application Services Command Reference*.

Using WAAS TCP Traceroute

The WAAS TCP Traceroute tool can help you troubleshoot network and connection issues, including asymmetric paths. You can use it to find a list of WAAS nodes between the client and server, and the configured and applied policies for a connection. From the Central Manager, you can choose any device in your WAAS network from which to run the traceroute.

To use the WAAS Central Manager TCP Traceroute tool, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Monitor > Troubleshoot > WAAS Tcptraceroute**.

Alternatively, you can choose a device first and then choose this menu item to run the traceroute from that device.

Step 2 From the WAAS Node drop-down list, choose a WAAS device from which to run the traceroute. (This item does not appear if you are in the device context.)

Step 3 In the Destination IP and Destination Port fields, enter the IP address and port of the destination to which you want to run the traceroute

Step 4 Click **Run TCPTraceroute** to display the results.

WAAS nodes in the traced path are displayed in the table below the fields. Use the filter settings in the Show drop-down list to filter the devices as needed. You can use a quick filter to filter on any value or show all devices.

You can display traceroute information from the CLI by using the **waas-tcptrace** EXEC command.

Another troubleshooting tool that you can use to trace connections on an ANC is the Connection Trace tool. For details, see the "AppNav Connection Tracing" section on page 4-37.

CHAPTER 18

Configuring SNMP Monitoring

This chapter describes how to configure SNMP traps, recipients, community strings and group associations, user security model groups, and user access permissions.

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network. The term WAE refers to WAE appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

This chapter contains the following sections:

About SNMP

Simple Network Management Protocol (SNMP) is an interoperable standards-based protocol that allows for external monitoring of WAAS devices through an SNMP agent.

An SNMP-managed network consists of the following primary components:

- Managed device—A network node that contains an SNMP agent and resides on a managed network. Managed devices include routers, access servers, switches, bridges, hubs, computer hosts, and printers. Each WAAS device running the WAAS software has an SNMP agent.

- SNMP agent—A software module that resides on a managed device. An agent has local knowledge of management information and translates that information into a form compatible with SNMP. The SNMP agent gathers data from the Management Information Base (MIB), which is the repository for information about device parameters and network data. The agent can also send traps, or notification of certain events, to the management system.
- Management station—Also known as the SNMP host, the management station uses SNMP to send the agent an SNMP Get request to obtain information from the WAAS device. The managed devices then collect and store management information and use SNMP to make this information available to the management station.

Before you can access this SNMP information, you must have deployed an SNMP management application on a management station. This SNMP management station is referred to as the SNMP host because it uses SNMP to send the device agent an SNMP Get request to obtain information from the WAAS device.

This section contains the following topics:

SNMP Communication Process

The SNMP management station and the SNMP agent that resides on a WAAS device use SNMP to communicate as follows:

1. The SNMP management station (the SNMP host) uses SNMP to request information from the WAAS device.
2. After receiving these SNMP requests, the SNMP agent on the WAAS device accesses a table that contains information about the individual device. This table, or database, is called a Management Information Base (MIB).

   Note The SNMP agent on the WAAS device only initiates communication with the SNMP host under unusual conditions; it will initiate communication when it has a trap it needs to send to the host. For more information on this topic, see the "Enabling SNMP Traps" section on page 18-13.

3. After locating the specified information in the MIB, the agent uses SNMP to send the information to the SNMP management station.

Figure 18-1 illustrates these SNMP operations for an individual WAAS device.

Figure 18-1 SNMP Components in a WAAS Network

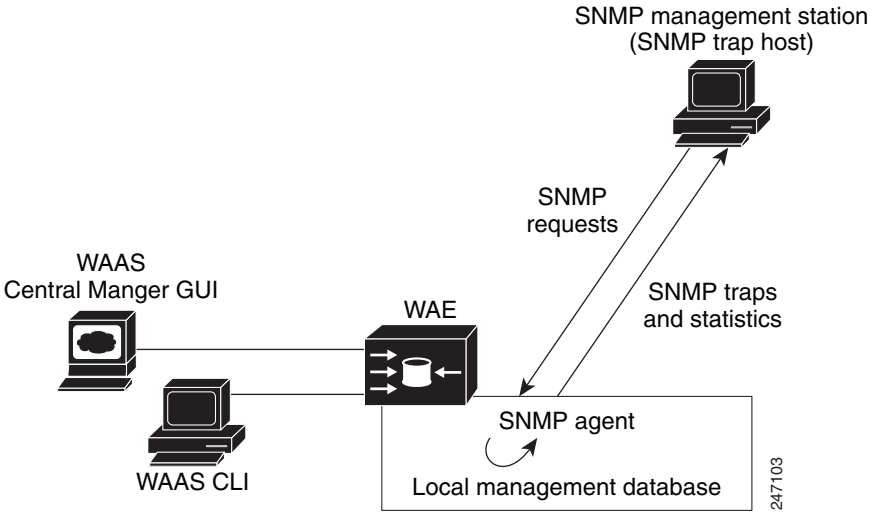

Supported SNMP Versions

The WAAS software supports the following versions of SNMP:

- Version 1 (SNMPv1)—This is the initial implementation of SNMP. See the RFC 1157 for a full description of its functionality.
- Version 2 (SNMPv2c)—This is the second release of SNMP, described in RFC 1902. It provides additions to data types, counter size, and protocol operations.
- Version 3 (SNMPv3)—This is the most recent version of SNMP, defined in RFC 2271 through RFC 2275.

Each Cisco device running WAAS software contains the software necessary to communicate information about device configuration and activity using the SNMP protocol.

SNMP Security Models and Security Levels

SNMPv1 and SNMPv2c do not have any security (that is, authentication or privacy) features to keep SNMP packet traffic confidential. As a result, packets on the wire can be detected and SNMP community strings compromised.

To solve the security shortcomings of SNMPv1 and SNMPv2c, SNMPv3 provides secure access to WAAS devices by authenticating and encrypting packets over the network. The SNMP agent in the WAAS software supports SNMPv3 as well as SNMPv1 and SNMPv2c.

The following security features are provided in SNMPv3:

- Message integrity—Ensures that nothing has interfered with a packet during transmission.
- Authentication—Determines that the message is from a valid source.
- Encryption—Scrambles the contents of a packet to prevent it from being seen by an unauthorized source.

SNMPv3 provides security models as well as security levels. A security model is an authentication process that is set up for a user and the group in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security process is used when an SNMP packet is handled. Three security models are available: SNMPv1, SNMPv2c, and SNMPv3.

Table 18-1 describes the combinations of security models and security levels.

Table 18-1 SNMP Security Models and Security Levels

Model	Level	Authentication	Encryption	Process
v1	noAuthNoPriv	Community string	No	Uses a community string match for user authentication.
v2c	noAuthNoPriv	Community string	No	Uses a community string match for user authentication.
v3	noAuthNoPriv	Username	No	Uses a username match for user authentication.
v3	AuthNoPriv	Message Digest 5 (MD5) or Secure Hash Algorithm (SHA)	No	Provides authentication based on the Hash-Based Message Authentication Code (HMAC)-MD5 or HMAC-SHA algorithms.
v3	AuthPriv	MD5 or SHA	Yes	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms. Provides Data Encryption Standard (DES) 56-bit encryption (packet authentication) based on the cipher block chaining (CBC)-DES (DES-56) standard.

The SNMPv3 agent can be used in the following modes:

- noAuthNoPriv mode (that is, no security mechanisms turned on for packets)
- AuthNoPriv mode (for packets that do not need to be encrypted using the privacy algorithm [DES 56])
- AuthPriv mode (for packets that must be encrypted; privacy requires that authentication be performed on the packet)

Using SNMPv3, users can securely collect management information from their SNMP agents without worrying that the data has been tampered with. Also, confidential information, such as SNMP set packets that change a Content Engine's configuration, can be encrypted to prevent their contents from being exposed on the wire. Also, the group-based administrative model allows different users to access the same SNMP agent with varying access privileges.

Supported MIBs

This section describes the Cisco-specific MIBs that are supported by WAAS. MIBs are listed in alphabetical order. The following Cisco-specific MIBs are supported:

- CISCO-APPNAV-MIB
- CISCO-CONTENT-ENGINE-MIB
- CISCO-ENTITY-ASSET-MIB
- CISCO-SMI
- CISCO-WAN-OPTIMIZATION-MIB
- ENTITY-MIB

- EVENT-MIB
- HOST-RESOURCES-MIB
- IF-MIB
- MIB-II
- SNMP-COMMUNITY-MIB
- SNMP-FRAMEWORK-MIB
- SNMP-NOTIFICATION-MIB
- SNMP-TARGET-MIB
- SNMP-USM-MIB
- SNMPv2-MIB
- SNMP-VACM-MIB

CISCO-APPNAV-MIB

This MIB provides information about AppNav objects. The following service context objects are supported when the WAAS device is in AppNav Controller mode:

- cAppNavServContextIndex
- cAppNavServContextName
- cAppNavServContextCurrOpState
- cAppNavServContextLastOpState
- cAppNavServContextIRState
- cAppNavServContextJoinState

The following AppNav controller group objects are supported:

- cAppNavACGIndex
- cAppNavACGName
- cAppNavACGServContextName

The following WAAS node group objects are supported:

- cAppNavSNGIndex
- cAppNavSNGName
- cAppNavSNGServContextName

The following AppNav controller objects are supported:

- cAppNavACIndex
- cAppNavACIpAddrType
- cAppNavACIpAddr
- cAppNavACServContextName
- cAppNavACACGName
- cAppNavACCurrentCMState

The following WAAS node objects are supported:

- cAppNavSNIndex

- cAppNavSNIpAddrType
- cAppNavSNIpAddr
- cAppNavSNServContextName
- cAppNavSNSNGName
- cAppNavSNCurrentCMState

CISCO-CDP-MIB

This MIB displays the ifIndex value of the local interface. For 802.3 repeaters on which the repeater ports do not have ifIndex values assigned, this value is a unique value for the port and is greater than any ifIndex value supported by the repeater. In this example, the specific port is indicated by the corresponding values of cdpInterfaceGroup and cdpInterfacePort, where these values correspond to the group number and the port number values of RFC 1516.

CISCO-CONFIG-MAN-MIB

This MIB represents a model of configuration data that exists in various locations:

- running—In use by the running system
- terminal—Saved to whatever hardware is attached as the terminal
- local—Saved locally in NVRAM or in flash memory
- remote—Saved to a server on the network

This MIB includes only operations that are specifically related to configuration, although some of the system functions can be used for general file storage and transfer.

CISCO-CONTENT-ENGINE-MIB

This is the MIB module for the Cisco WAE device from Cisco Systems, Inc. The following objects from this MIB are supported:

- cceAlarmCriticalCount
- cceAlarmMajorCount
- cceAlarmMinorCount
- cceAlarmHistTable

CISCO-ENTITY-ASSET-MIB

This MIB monitors the asset information of items in the ENTITY-MIB (RFC 2037) entPhysicalTable. This MIB lists the orderable part number, serial number, hardware revision, manufacturing assembly number and revision, firmware ID and revision (if any) and software ID and revision (if any) of relevant entities listed in ENTITY-MIB entPhysicalTable.

Entities that have none of this data available are not listed in this MIB. The table in this MIB is sparsely populated, so some variables may not exist for a particular entity at a particular time. For example, a row that represents a powered-off module may have no values for software ID (ceAssetSoftwareID) and revision (ceAssetSoftwareRevision). Similarly, a power supply would probably never have firmware or software information listed in the table.

Although the data may have other items encoded in it (for example, a manufacturing date in the serial number), consider all data items to be a single unit. Do not decompose the items or parse them. Use only string equal and unequal operations on them.

CISCO-SMI

This is the MIB module for Cisco Enterprise Structure of Management Information. There is nothing to query in this MIB; it describes the structure of Cisco MIBs.

CISCO-WAN-OPTIMIZATION-MIB

This MIB provides information about the status and statistics associated with optimization and the application accelerators.

The following TFO statistics objects are supported:

- cwoTfoStatsTotalOptConn
- cwoTfoStatsActiveOptConn
- cwoTfoStatsMaxActiveConn
- cwoTfoStatsActiveOptTCPPlusConn
- cwoTfoStatsActiveOptTCPOnlyConn
- cwoTfoStatsActiveOptTCPPrepConn
- cwoTfoStatsActiveADConn
- cwoTfoStatsReservedConn
- cwoTfoStatsPendingConn
- cwoTfoStatsActivePTConn
- cwoTfoStatsTotalNormalClosedConn
- cwoTfoStatsResetConn
- cwoTfoStatsLoadStatus

The following general application accelerator statistics objects are supported:

- cwoAoStatsName
- cwoAoStatsIsConfigured
- cwoAoStatsIsLicensed
- cwoAoStatsOperationalState
- cwoAoStatsStartUpTime
- cwoAoStatsLastResetTime
- cwoAoStatsTotalHandledConn
- cwoAoStatsTotalOptConn
- cwoAoStatsTotalHandedOffConn
- cwoAoStatsTotalDroppedConn
- cwoAoStatsActiveOptConn
- cwoAoStatsPendingConn
- cwoAoStatsMaxActiveOptConn
- cwoAoStatsLoadStatus
- cwoAoStatsBwOpt

The following SMB application accelerator statistics objects are supported:

- cwoAoSmbxStatsBytesReadCache
- cwoAoSmbxStatsBytesWriteCache
- cwoAoSmbxStatsBytesReadServer
- cwoAoSmbxStatsBytesWriteServer
- cwoAoSmbxStatsBytesReadClient
- cwoAoSmbxStatsBytesWriteClient
- cwoAoSmbxStatsProcessedReqs
- cwoAoSmbxStatsActiveReqs
- cwoAoSmbxStatsTotalTimedOutReqs
- cwoAoSmbxStatsTotalRemoteReqs
- cwoAoSmbxStatsTotalLocalReqs
- cwoAoSmbxStatsRemoteAvgTime
- cwoAoSmbxStatsLocalAvgTime
- cwoAoSmbxStatsRACacheHitCount
- cwoAoSmbxStatsMDCacheHitCount
- cwoAoSmbxStatsRACacheHitRate
- cwoAoSmbxStatsMDCacheHitRate
- cwoAoSmbxStatsMaxRACacheSize
- cwoAoSmbxStatsMaxMDCacheSize
- cwoAoSmbxStatsMDCacheSize
- cwoAoSmbxStatsRAEvictedAge
- cwoAoSmbxStatsRTT
- cwoAoSmbxStatsTotalRespTimeSaving
- cwoAoSmbxStatsOpenFiles
- cwoAoSmbxStatsTotalFilesInRACache

The following HTTP application accelerator statistics objects are supported:

- cwoAoHttpxStatsTotalSavedTime
- cwoAoHttpxStatsTotalRTT
- cwoAoHttpxStatsTotalMDCMTime
- cwoAoHttpxStatsEstSavedTime

The following MAPI application accelerator statistics objects are supported:

- cwoAoMapixStatsUnEncrALRT
- cwoAoMapixStatsUnEncrARRT
- cwoAoMapixStatsTotalUnEncrLRs
- cwoAoMapixStatsTotalUnEncrRRs
- cwoAoMapixStatsUnEncrAvgRedTime
- cwoAoMapixStatsEncrALRT

- cwoAoMapixStatsEncrARRT
- cwoAoMapixStatsTotalEncrLRs
- cwoAoMapixStatsTotalEncrRRs
- cwoAoMapixStatsEncrAvgRedTime

The following NFS application accelerator statistics objects are supported:

- cwoAoNfsxStatsALRT
- cwoAoNfsxStatsARRT
- cwoAoNfsxStatsTotalLRs
- cwoAoNfsxStatsTotalRRs
- cwoAoNfsxStatsEstTimeSaved

The following video application accelerator statistics objects are supported:

- cwoAoVideoxStatsTotalInBytes
- cwoAoVideoxStatsTotalOutBytes

The following application statistics objects are supported:

- cwoAppStatsAppName
- cwoAppStatsOriginalBytes
- cwoAppStatsOptimizedBytes
- cwoAppStatsPTBytes

The following optimization policy map statistics objects are supported:

- cwoPmapStatsType
- cwoPmapStatsName
- cwoPmapStatsDescr
- cwoPmapStatsTotalConns
- cwoPmapStatsTotalBytes
- cwoPmapStatsTotalPTConns
- cwoPmapStatsTotalPTBytes

The following optimization class map statistics objects are supported:

- cwoCmapStatsType
- cwoCmapStatsName
- cwoCmapStatsDescr
- cwoCmapStatsTotalConns
- cwoCmapStatsTotalBytes
- cwoCmapStatsTotalPTConns
- cwoCmapStatsTotalPTBytes

ENTITY-MIB

This is the MIB module for representing multiple logical entities supported by a single SNMP agent. This MIB is documented in RFC 2737. The following groups from this MIB are supported:

- entityPhysicalGroup
- entityLogicalGroup

The entConfigChange notification is supported.

EVENT-MIB

This MIB defines event triggers and actions for network management purposes. The MIB is published as RFC 2981.

HOST-RESOURCES-MIB

This MIB manages host systems. The term “host” implies any computer that communicates with other similar computers connected to the Internet. The HOST-RESOURCES-MIB does not necessarily apply to devices whose primary function is communications services (terminal servers, routers, bridges, monitoring equipment). This MIB provides attributes that are common to all Internet hosts, for example, personal computers and systems that run variants of UNIX. The following objects from this MIB are not supported:

- HrPrinterEntry
- hrSWOSIndex
- hrSWInstalledGroup

IF-MIB

This MIB supports querying for interface-related statistics including 64-bit interface counters. These counters include received and sent octets, unicast, multicast, and broadcast packets on the device interfaces. All the objects from ifXEntry are supported except for ifCounterDiscontinuityTime. This MIB is documented in RFC 2233.

Loopback interface and virtual blade interface information are not reported.

MIB-II

MIB-II is the Internet Standard MIB. The MIB-II is documented in RFC 1213 and is for use with network management protocols in TCP/IP-based internets. This MIB is found in the RFC1213-MIB file in the v1 directory on the download site (other MIBs are in the v2 directory). The following objects from this MIB are not supported:

- ifInUnknownProtos
- ifOutNUcastPkts
- ipRouteAge
- TcpConnEntry group
- egpInMsgs
- egpInErrors
- egpOutMsgs
- egpOutErrors
- EgpNeighEntry group
- egpAs

SNMP-COMMUNITY-MIB

This MIB is documented in RFC 2576.

SNMP-FRAMEWORK-MIB

This MIB is documented in RFC 2571.

SNMP-NOTIFICATION-MIB

This MIB is documented in RFC 3413.

SNMP-TARGET-MIB

This MIB is documented in RFC 3413.

SNMP-USM-MIB

This MIB is documented in RFC 2574.

SNMPv2-MIB

This MIB is documented in RFC 1907. WAAS supports the following notifications from this MIB:

- coldStart
- linkUp
- linkDown
- authenticationFailure

SNMP-VACM-MIB

This MIB is documented in RFC 2575.

Downloading MIB Files

You can download the MIB files for most of the MIBS that are supported by a device that is running the WAAS software from the following Cisco FTP site:

ftp://ftp.cisco.com/pub/mibs/v2

You can download the RFC1213-MIB file (for MIB-II) from the following Cisco FTP site:

ftp://ftp.cisco.com/pub/mibs/v1

The MIB objects that are defined in each MIB are described in the MIB files at the above FTP sites and are self-explanatory.

Enabling the SNMP Agent on a WAAS Device

By default, the SNMP agent on WAAS devices is disabled and an SNMP community string is not defined. The SNMP community string is used as a password for authentication when accessing the SNMP agent on a WAAS device. To be authenticated, the Community Name field of any SNMP message sent to the WAAS device must match the SNMP community string defined on the WAAS device.

The SNMP agent on a WAAS device is enabled when you define the SNMP community string on the device. The WAAS Central Manager GUI allows you to define the SNMP community string on a device or device group.

If the SNMPv3 protocol is going to be used for SNMP requests, the next step is to define an SNMP user account that can be used to access a WAAS device through SNMP. For more information on how to create an SNMPv3 user account on a WAAS device, see the "Creating an SNMP User" section on page 18-22.

Checklist for Configuring SNMP

Table 18-2 describes the process for enabling SNMP monitoring on a WAAS device or device group.

Table 18-2 Checklist for Configuring SNMP

Task	Additional Information and Instructions
1. Prepare for SNMP monitoring.	For more information, see the "Preparing for SNMP Monitoring" section on page 18-13.
2. Select the SNMP traps that you want to enable.	The WAAS Central Manager provides a wide-range of traps that you can enable on a WAAS device or device group. For more information, see the "Enabling SNMP Traps" section on page 18-13. To define additional traps, see the "Defining SNMP Traps" section on page 18-16.
3. Specify the SNMP host that receives the SNMP traps.	Specify the SNMP host to that the WAAS device or device group should send their traps to. You can specify multiple hosts so different WAAS devices send traps to different hosts. For more information, see the "Specifying the SNMP Host" section on page 18-18.
4. Specify the SNMP community string.	Specify the SNMP community string so external users can read or write to the MIB. For more information, see the "Specifying the SNMP Community String" section on page 18-19.
5. Set up SNMP views.	To restrict an SNMP group to a specific view, you must create a view that specifies the MIB subtree that you want the group to view. For more information, see the "Creating SNMP Views" section on page 18-20.
6. Create an SNMP group.	You must set up an SNMP group if are going to create any SNMP users or want to restrict a group to view a specific MIB subtree. For more information, see the "Creating an SNMP Group" section on page 18-21.
7. Create an SNMP user.	If the SNMPv3 protocol is going to be used for SNMP requests, you must create at least one SNMPv3 user account on the WAAS device in order for the WAAS device to be accessed through SNMP. For more information, see the "Creating an SNMP User" section on page 18-22.
8. Configure SNMP contact settings.	For more information, see the "Configuring SNMP Contact Settings" section on page 18-24.

Preparing for SNMP Monitoring

Before you configure your WAAS network for SNMP monitoring, complete the following preparation tasks:

- Set up the SNMP host (management station) that the WAAS devices will use to send SNMP traps.
- Determine if all your WAAS devices will be sending traps to the same host, or to different hosts. Write down the IP address or hostname of each SNMP host.
- Obtain the community string used to access the SNMP agents.
- Determine if you want to create SNMP groups so you can restrict views by group.
- Determine what additional SNMP traps you need.
- Clock synchronization between the devices in a WAAS network is important. On each WAAS device, be sure to set up a Network Time Protocol (NTP) server to keep the clocks synchronized.

Enabling SNMP Traps

To enable a WAAS device to send SNMP traps, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Monitoring > SNMP > General Settings**. The SNMP General Settings window appears. (See Figure 18-2.) Table 18-3 describes the fields in this window.

Figure 18-2 SNMP General Settings Window

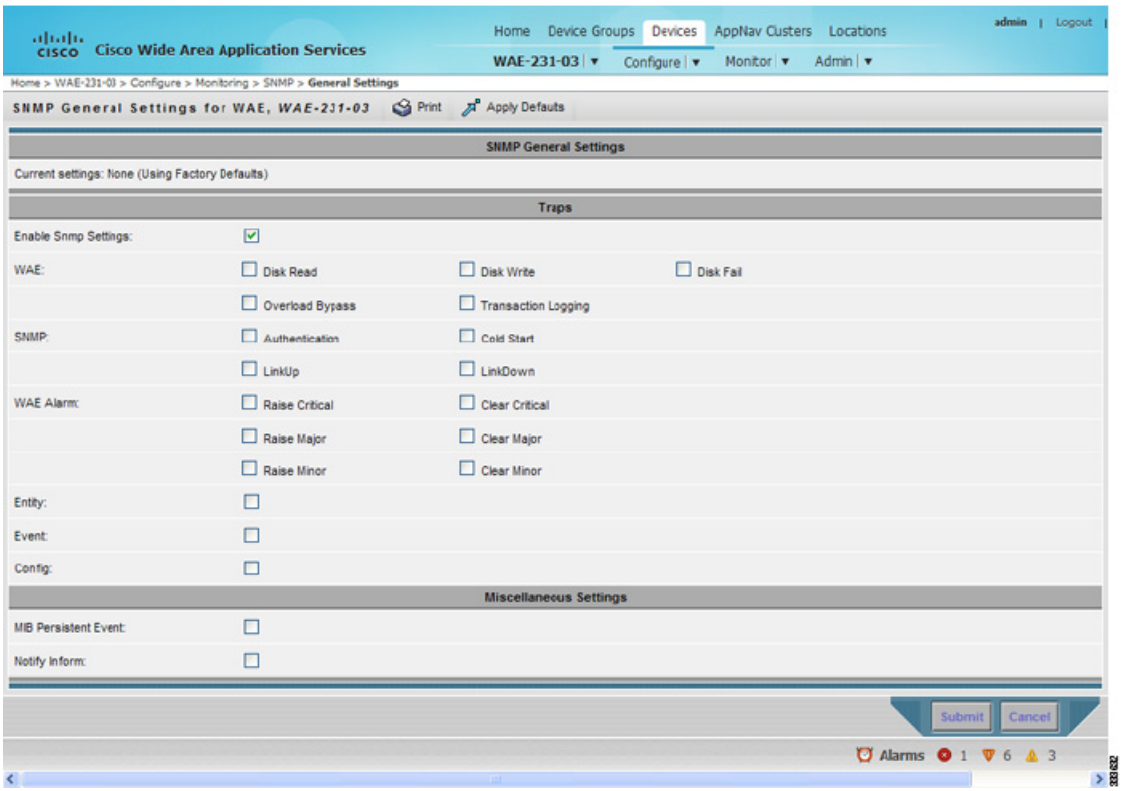

Table 18-3 SNMP General Settings

GUI Parameter	Function
Traps	
Enable Snmp Settings	Enables SNMP traps.
WAE	Enables SNMP WAE traps: • Disk Read—Enables disk read error trap. • Disk Write—Enables disk write error trap. • Disk Fail—Enables disk failure error trap. • Overload Bypass—Enables WCCP overload bypass error trap. • Transaction Logging—Enables transaction log write error trap.
SNMP	Enables SNMP-specific traps: • Authentication—Enables authentication trap. • Cold Start—Enables cold start trap. • LinkUp—Link up trap. • LinkDown—Link down trap.

Table 18-3 SNMP General Settings (continued)

GUI Parameter	Function
WAE Alarm	Enables WAE alarm traps: • Raise Critical—Enables raise-critical alarm trap • Clear Critical—Enables clear-critical alarm trap • Raise Major—Enables raise-major alarm trap • Clear Major—Enables clear-major alarm trap • Raise Minor—Enables raise-minor alarm trap • Clear Minor—Enables clear-minor alarm trap
Entity	Enables SNMP entity traps.
Event	Enables the Event MIB.
Config	Enables CiscoConfigManEvent error traps.
Miscellaneous Settings	
MIB Persistent Event	Enables persistence for the SNMP Event MIB. (This check box is not shown when the selected device is a Central Manager.)
Notify Inform	Enables the SNMP notify inform request. Inform requests are more reliable than traps but consume more resources in the router and in the network. Traps are unreliable because the receiver does not send any acknowledgment when it receives a trap. The sender cannot determine if the trap was received. However, an SNMP manager that receives an inform request acknowledges the message with an SNMP response. If the sender never receives a response, the inform request can be sent again. Thus, informs are more likely to reach their intended destinations.

Step 3 Check the appropriate check boxes to enable SNMP traps.

Step 4 Click **Submit**.

A "Click Submit to Save" message appears in red next to the current settings when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The Reset button is visible only when you apply default or device group settings to change the current device settings but the settings have not yet been submitted.

To enable SNMP traps from the CLI, you can use the **snmp-server enable traps** global configuration command.

To control access to the SNMP agent by an external SNMP server, use the **snmp-server access-list** global configuration command to apply an SNMP ACL.

Note If you are using an SNMP server ACL, you must permit the loopback interface.

Note If you override the device group settings from the SNMP General Settings window, the Central Manager deletes the SNMP community, SNMP group, SNMP user, SNMP view, and SNMP host settings. You are asked to confirm this behavior.

To define additional SNMP traps for other MIB objects of interest to your particular configuration, see the "Defining SNMP Traps" section on page 18-16.

Defining SNMP Traps

To define additional SNMP traps for other MIB objects of interest to your particular configuration, follow these steps to create additional SNMP triggers:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Monitoring > SNMP > Trigger**. The SNMP Trigger List Entries window appears. The columns in this window are the same as the parameters described in Table 18-4.

Step 3 In the taskbar, click the **Create New SNMP Trigger List Entry** icon. The Creating New SNMP Trigger window appears. Table 18-4 describes the fields in this window.

Table 18-4 Creating New SNMP Trigger Settings

GUI Parameter	Function
MIB Name	MIB variable name of the object that you want to monitor.
Wild Card	(Optional) Check this check box if the MIB Name value is a wildcard. Note that this check box is disabled when editing the SNMP Trigger.
Frequency	Number of seconds (60–600) to wait between trigger samples.

Table 18-4 Creating New SNMP Trigger Settings (continued)

GUI Parameter	Function
Test	Test used to trigger the SNMP trap. Choose one of the following tests: • absent—A specified MIB object that was present at the last sampling is no longer present as of the current sampling. • equal—The value of the specified MIB object is equal to the specified threshold. • falling—The value of the specified MIB object has fallen below the specified threshold value. After a trap is generated against this condition, another trap for this same condition is not generated until the sampled MIB object value rises above the threshold value and then falls below the falling threshold value again. • greater-than—The value of the specified MIB object is greater than the specified threshold value. • less-than—The value of the specified MIB object is less than the specified threshold value. • on-change—The value of the specified MIB object has changed since the last sampling. • present—A specified MIB object is present as of the current sampling that was not present at the previous sampling. • rising—The value of the specified MIB object has risen above the specified threshold. After a trap is generated against this condition, another trap for this same condition is not generated until the sampled MIB object value falls below the threshold value and then rises above the rising threshold value again.
Sample Type	(Optional) Sample type, as follows: • absolute—The test is evaluated against a fixed integer value between zero and 2147483647. • delta—The test is evaluated against the change in the MIB object value between the current sampling and the previous sampling.
Threshold Value	Threshold value of the MIB object. This field is not used if absent, on-change, or present is chosen in the Test drop-down list.
MIB Var1 MIB Var2 MIB Var3	(Optional) Names of up to three alternate MIB variables to add to the notification. Validation of these names is not supported, so be sure to enter them correctly.
Comments	Description of the trap.

Step 4 In the appropriate fields, enter the MIB name, frequency, test, sample type, threshold value, and comments.

Note You can create valid triggers only on read-write and read-only MIB objects. If you create a trigger on a read-create MIB object, it is deleted from the Central Manager configuration after one one data feed poll cycle.

Step 5 Click **Submit**.

The new SNMP trigger is listed in the SNMP Trigger List window.

You can edit an SNMP trigger by clicking the **Edit** icon next to the MIB name in the SNMP Trigger List Entries window.

You can delete an SNMP trigger by clicking the **Edit** icon next to the MIB name and then clicking the **Delete** taskbar icon.

Note If you delete any of the default SNMP triggers, they will be restored after a reload.

You can use the **snmp trigger** EXEC command to define SNMP traps from the CLI.

To control access to the SNMP agent by an external SNMP server, use the **snmp-server access-list** global configuration command to apply an SNMP ACL.

Note If you are using an SNMP server ACL, you must permit the loopback interface.

Aggregating SNMP Triggers

An individual WAE device can have custom SNMP triggers defined and can belong to device groups that have other custom SNMP triggers defined.

In the SNMP Trigger List Entries window, the Aggregate Settings radio button controls how SNMP triggers are aggregated for an individual device, as follows:

- Choose **Yes** if you want to configure the device with all custom SNMP triggers that are defined for itself and for device groups to which it belongs.
- Choose **No** if you want to limit the device to just the custom SNMP triggers that are defined for itself.

When you change the setting, you get the following confirmation message: “This option will take effect immediately and will affect the device configuration. Do you wish to continue?” Click **OK** to continue.

Specifying the SNMP Host

Hosts are listed in the order in which they have been created. The maximum number of SNMP hosts that can be created is four.

To specify the SNMP host, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Monitoring > SNMP > Host**. The SNMP Hosts window appears.

Step 3 In the taskbar, click the **Create New SNMP Host** icon. The Creating New SNMP Host window appears. Table 18-5 describes the fields in this window.

Table 18-5 SNMP Host Settings

GUI Parameter	Function
Trap Host	Hostname or IP address of the SNMP trap host that is sent in SNMP trap messages from the WAE. This is a required field.
Community/User	Name of the SNMP community or user (64 characters maximum) that is sent in SNMP trap messages from the WAE. This is a required field.
Authentication	Security model to use for sending notification to the recipient of an SNMP trap operation. Choose one of the following options from the drop-down list: • **No-auth**—Sends notification without any security mechanism. • **v2c**—Sends notification using Version 2c security. • **v3-auth**—Sends notification using SNMP Version 3 AuthNoPriv. • **v3-noauth**—Sends notification using SNMP Version 3 NoAuthNoPriv security. • **v3-priv**—Sends notification using SNMP Version 3 AuthPriv security.
Retry	Number of retries (1–10) allowed for the inform request. The default is 2 tries.
Timeout	Timeout for the inform request in seconds (1–1000). The default is 15 seconds.

Step 4 Enter the hostname or IP address of an SNMP trap host, SNMP community or user name, security model to send notification, and retry count and timeout for inform requests.

Step 5 Click **Submit**.

To specify the SNMP host from the CLI, you can use the **snmp-server host** global configuration command.

Specifying the SNMP Community String

An SNMP community string is the password used to access an SNMP agent that resides on WAAS devices. There are two types of community strings: group and read-write. Community strings enhance the security of your SNMP messages.

Community strings are listed in the order in which they have been created. The maximum number of SNMP communities that can be created is ten. By default, an SNMP agent is disabled, and a community string is not configured. When a community string is configured, it permits read-only access to all agents by default.

To enable the SNMP agent and configure a community string to permit access to the SNMP agent, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Monitoring > SNMP > Community**. The SNMP Community Strings window appears.

Step 3 In the taskbar, click the **Create New SNMP Community String** icon. The Creating New SNMP Community String window appears. Table 18-6 describes the fields in this window.

Table 18-6 SNMP Community Settings

GUI Parameter	Function
Community	Community string used as a password for authentication when you access the SNMP agent of the WAE. The "Community Name" field of any SNMP message sent to the WAE must match the community string defined here in order to be authenticated. Entering a community string enables the SNMP agent on the WAE. You can enter a maximum of 64 characters in this field. This is a required field.
Group name/rw	Group to which the community string belongs. The Read/Write option allows a read or write group to be associated with this community string. The Read/Write option permits access to only a portion of the MIB subtree. Choose one of the following three options from the drop-down list: • **None**—Choose this option if you do not want to specify a group name to be associated with the community string. The Group Name field remains disabled if you select this option. • **Group**—Choose this option if you want to specify a group name. • **Read/Write**—Choose this option if you want to allow read-write access to the group associated with a community string. The Group Name field remains disabled if you select this option. This is a required field.
Group Name	Name of the group to which the community string belongs. You can enter a maximum of 64 characters in this field. This field is available only if you have chosen the Group option in the previous field.

Step 4 In the appropriate fields, enter the community string, choose whether or not read-write access to the group is allowed, and enter the group name.

Step 5 Click **Submit**.

To configure a community string from the CLI, you can use the **snmp-server community** global configuration command.

Creating SNMP Views

To restrict a group of users to view a specific MIB tree, you must create an SNMP view using the WAAS Central Manager GUI. Once you create the view, you need to create an SNMP group and SNMP users that belong to this group as described in later sections.

Views are listed in the order in which they have been created. The maximum number of views that can be created is ten.

To create a Version 2 SNMP (SNMPv2) MIB view, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Monitoring > SNMP > View**. The SNMP Views window appears.

Step 3 In the taskbar, click the **Create New View** icon. The Creating New SNMP View window appears. Table 18-7 describes the fields in this window.

Table 18-7 SNMPv2 View Settings

GUI Parameter	Function
Name	String representing the name of this family of view subtrees (64 characters maximum). The family name must be a valid MIB name such as ENTITY-MIB. This is a required field.
Family	Object identifier (64 characters maximum) that identifies a subtree of the MIB. This is a required field.
View Type	View option that determines the inclusion or exclusion of the MIB family from the view. Choose one of the following two options from the drop-down list: • **Included**—The MIB family is included in the view. • **Excluded**—The MIB family is excluded from the view.

Step 4 In the appropriate fields, enter the view name, the family name, and the view type.

Step 5 Click **Submit**.

Step 6 Create an SNMP group that will be assigned to this view as described in the section that follows.

To create an SNMP view from the CLI, you can use the **snmp-server view** global configuration command.

Creating an SNMP Group

You must set up an SNMP group if you are going to create any SNMP users or want to restrict a group of users to view a specific MIB subtree.

Groups are listed in the order in which they have been created. The maximum number of SNMP groups that can be created is ten.

To define a user security model group, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Monitoring > SNMP > Group**. The SNMP Group Strings for WAE window appears.

Step 3 In the taskbar, click the **Create New SNMP Group String** icon. The Creating New SNMP Group String for WAE window appears. Table 18-8 describes the fields in this window.

Table 18-8 SNMP Group Settings

GUI Parameter	Function
Name	Name of the SNMP group. You can enter a maximum of 64 characters. This is a required field.
Sec Model	Security model for the group. Choose one of the following options from the drop-down list: • **v1**—Version 1 security model (SNMP Version 1 [noAuthNoPriv]). • **v2c**—Version 2c security model (SNMP Version 2 [noAuthNoPriv]). • **v3-auth**—User security level SNMP Version 3 AuthNoPriv. • **v3-noauth**—User security level SNMP Version 3 noAuthNoPriv. • **v3-priv**—User security level SNMP Version 3 AuthPriv. **Note** A group defined with the SNMPv1 or SNMPv2c security model should not be associated with SNMP users; they should only be associated with the community strings.
Read View	Name of the view (a maximum of 64 characters) that enables you only to view the contents of the agent. By default, no view is defined. In order to provide read access to users of the group, a view must be specified. For information on creating SNMP views, see the "Creating SNMP Views" section on page 18-20.
Write View	Name of the view (a maximum of 64 characters) that enables you to enter data and configure the contents of the agent. By default, no view is defined. For information on creating SNMP views, see the "Creating SNMP Views" section on page 18-20.
Notify View	Name of the view (a maximum of 64 characters) that enables you to specify a notify, inform, or trap. By default, no view is defined. For information on creating SNMP views, see the "Creating SNMP Views" section on page 18-20.

Step 4 In the appropriate fields, enter the SNMP group configuration name, the security model, and the names of the read, write, and notify views.

Step 5 Click **Submit**.

Step 6 Create SNMP users that belong to this new group as described in the section that follows.

To create an SNMP group from the CLI, you can use the **snmp-server group** global configuration command.

Creating an SNMP User

Users are listed in the order in which they have been created. The maximum number of users that can be created is ten.

To define a user who can access the SNMP engine, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Monitoring > SNMP > User**. A list of SNMP users for the device or device group appears.

Step 3 In the taskbar, click the **Create New SNMP User** icon. The Creating New SNMP User window appears. Table 18-9 describes the fields in this window.

Table 18-9 SNMP User Settings

GUI Parameter	Function
Name	String representing the name of the user (32 characters maximum) who can access the device or device group. This is a required field.
Group	Name of the group (64 characters maximum) to which the user belongs. This is a required field.
Remote SNMP ID	Globally unique identifier for a remote SNMP entity (10 to 64 characters). To send an SNMPv3 message to the WAE, at least one user with a remote SNMP ID must be configured on the WAE. The SNMP ID must be entered in octet string format. Only hexadecimal characters and the colon (:) are allowed in this field. If any colons appear in the entered string, they are removed when the page is submitted.
Authentication Algorithm	Authentication algorithm that ensures the integrity of SNMP packets during transmission. Choose one of the following three options from the drop-down list: • **No-auth**—Requires no security mechanism to be turned on for SNMP packets. • **MD5**—Provides authentication based on the hash-based Message Authentication Code Message Digest 5 (HMAC-MD5) algorithm. • **SHA**—Provides authentication based on the hash-based Message Authentication Code Secure Hash (HMAC-SHA) algorithm.
Authentication Password	String (256 characters maximum) that configures the user authentication (HMAC-MD5 or HMAC-SHA) password. The number of characters is adjusted to fit the display area if it exceeds the limit for display. The following special characters are not supported: space, backwards single quote (`), single quote (’), double quote (”), pipe (|), or question mark (?). This field is optional if the **no-auth** option is chosen for the authentication algorithm. Otherwise, this field must contain a value.
Confirmation Password	Authentication password for confirmation. The reentered password must be the same as the one entered in the previous field.
Private Password	String (256 characters maximum) that configures the authentication (HMAC-MD5 or HMAC-SHA) parameters to enable the SNMP agent to receive packets from the SNMP host. The number of characters is adjusted to fit the display area if it exceeds the limit for display. The following special characters are not supported: space, backwards single quote (`), double quote (”), pipe (|), or question mark (?).
Confirmation Password	Private password for confirmation. The reentered password must be the same as the one entered in the previous field.

Step 4 In the appropriate fields, enter the username, the group to which the user belongs, the engine identity of the remote entity to which the user belongs, the authentication algorithm used to protect SNMP traffic from tampering, the user authentication parameters, and the authentication parameters for the packet.

Step 5 Click **Submit**.

To create an SNMP user from the CLI, you can use the **snmp-server user** global configuration command.

Configuring SNMP Asset Tag Settings

To configure SNMP asset tag settings, which create values in the CISCO-ENTITY-ASSET-MIB, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Monitoring > SNMP > Asset Tag**. The SNMP Asset Tag Settings window appears.

Step 3 In the Asset Tag Name field, enter a name for the asset tag.

Step 4 Click **Submit**.

To configure SNMP asset tag settings from the CLI, you can use the **asset tag** global configuration command.

Configuring SNMP Contact Settings

To configure SNMP contact settings, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name* (or **Device Groups >** *device-group-name*).

Step 2 Choose **Configure > Monitoring > SNMP > Contact Information**. The SNMP Contact Settings window appears.

Step 3 Enter a contact name and location in the provided fields.

Step 4 Click **Submit**.

To configure SNMP contact settings from the CLI, you can use the **snmp-server contact** global configuration command.

Configuring SNMP Trap Source Settings

To configure the source interface from which SNMP traps are sent, follow these steps:

Step 1 From the WAAS Central Manager menu, choose **Devices >** *device-name*. (This setting is not supported from device groups.)

Step 2 Choose **Configure > Monitoring > SNMP > Trap Source**. The SNMP Trap Source Settings window appears.

Step 3 From the Trap Source drop-down list, choose the interface to be used as the trap source. From the available physical, standby, and port-channel interfaces, only those with IP addresses are shown in the list. For vWAAS devices, virtual interfaces with assigned IP addresses are shown in the list.

Note An interface assigned as a trap source cannot be removed until it is unassigned as a trap source.

Step 4 Click **Submit**.

To configure SNMP trap source settings from the CLI, you can use the **snmp-server trap-source** global configuration command.

APPENDIX A

Predefined Optimization Policy

The WAAS software includes over 200 predefined optimization policy rules that help your WAAS system classify and optimize some of the most common traffic on your network. Table A-1 lists the predefined applications and class maps that WAAS will either optimize or pass through based on the policy rules that are provided with the system.

Before you create an optimization policy, we recommend that you review the predefined policy rules and modify them as appropriate. Often, you can more easily modify an existing policy rule than create a new one.

When reviewing Table A-1, note the following information:

- The subheadings represent the application names, and the associated class maps are listed under these subheadings. For example, Authentication is a type of application and Kerberos is a class map for that application.
- Applications and class maps with the word (*monitored*) next to them are monitored by the WAAS Central Manager, which can monitor statistics for up to 25 applications and 25 class maps at a time. To view statistics for one of the unmonitored applications, use one of the following methods:
  - Use the WAAS CLI, which can display statistics for all applications and class maps on a WAAS device. For more information, see the *Cisco Wide Area Application Services Command Reference*.
  - Modify the application or class map settings so the WAAS Central Manager GUI displays statistics for the desired application or class map. For more information, see Chapter 13, "Configuring Application Acceleration."
- WAAS Express devices have similar default policy rules but provide application acceleration only for HTTP, SSL, and SMB traffic. Where a different application accelerator is listed in Table A-1, it is not part of the WAAS Action for a WAAS Express device.

The WAAS software uses the following optimization technologies based on the type of traffic that it encounters:

- TFO (transport flow optimization)—A collection of optimization technologies such as automatic windows scaling, increased buffering, and selective acknowledgement that optimize all TCP traffic over your network.
- DRE (data redundancy elimination)—A compression technology that reduces the size of transmitted data by removing redundant information before sending the shortened data stream over the WAN. DRE operates on significantly larger streams and maintains a much larger compression history than LZ compression. DRE can use bidirectional, unidirectional, or adaptive caching. Unless noted in Table A-1, DRE caching is bidirectional.
- LZ (compression)—Another compression technology that operates on smaller data streams and keeps limited compression history compared to DRE.

- Application accelerator—A collection of individual application accelerators for the following traffic types: CIFS, EPM, HTTP, ICA, MAPI, NFS, SSL, and streaming video. (Some application accelerators are not available on WAAS Express devices.)

Table A-1 Predefined Traffic Policy Rules

Application/Class Map	WAAS Action	Destination Ports
class-default (*monitored*)	LZ+TFO+ DRE-adaptive	All ports not included in other class maps
Authentication		
apple-sasl	Passthrough	3659
auth	Passthrough	113
Kerberos	Passthrough	88, 888, 2053
kerberos-adm (*monitored*)	Passthrough	749
klogin	Passthrough	543
kpasswd	Passthrough	464
kshell	Passthrough	544
TACACS	Passthrough	49
tell	Passthrough	754
Backup *(monitored)*		
Amanda	TFO	10080
backup-express	TFO	6123
CommVault	TFO	8400–8403
connected	TFO	16384
IBM-TSM	LZ+TFO+ DRE-unidirectional	1500-1502
Legato-NetWorker	TFO	7937, 7938, 7939
Legato-RepliStor	TFO	7144, 7145
Veritas-BackupExec (*monitored*)	TFO	1125, 3527, 6101, 6102, 6106
Veritas-NetBackup	TFO	13720, 13721, 13782, 13785
CAD		
PDMWorks	LZ+TFO+DRE	30000, 40000
Call-Management		
Cisco-CallManager	Passthrough	2443, 2748
cisco-q931-backhaul	Passthrough	2428
cisco-sccp	Passthrough	2000–2002
h323hostcall	Passthrough	1720
h323hostcallsc	Passthrough	1300
mgcp-callagent	Passthrough	2727
mgcp-gateway	Passthrough	2427
sip	Passthrough	5060

Table A-1 Predefined Traffic Policy Rules (continued)

Application/Class Map	WAAS Action	Destination Ports
sip-tls	Passthrough	5061
VoIP-Control	Passthrough	1718, 1719, 11000–11999
CIFS		
CIFS *(monitored)*	LZ+TFO+DRE+ CIFS accelerator	139, 445
Citrix		
Citrix-ICA (*monitored*)	TFO+ ICA accelerator	1494
Citrix-CGP (*monitored*)	TFO+ ICA accelerator	2598
Conferencing		
cuseeme	Passthrough	7640, 7642, 7648, 7649
ezMeeting	Passthrough	10101–10103, 26260, 26261
MS-NetMeeting (*monitored*)	Passthrough	522, 1503, 1731
proshare	Passthrough	5713–5717
PSOM-MTLS	Passthrough	8057
VocalTec	Passthrough	1490, 6670, 25793, 22555
Console		
cmd	Passthrough	514
exec	Passthrough	512
login	Passthrough	513
sshell	Passthrough	614
Telnet	Passthrough	23, 107
Telnets	Passthrough	992
Content-Management *(monitored)*		
dmdocbroker	LZ+TFO+DRE	1489
Filenet	LZ+TFO+DRE	32768–32774
Directory-Services *(monitored)*		
LDAP	LZ+TFO+ DRE-unidirectional	389, 8404
ldaps	Passthrough	636
msft-gc	LZ+TFO+ DRE-unidirectional	3268
msft-gc-ssl	Passthrough	3269
Email-and-Messaging *(monitored)*		
ccmail	LZ+TFO+DRE	3264
groupwise	LZ+TFO+DRE	1677, 2800, 3800, 7100, 7101, 7180, 7181, 7205, 9850

Table A-1 Predefined Traffic Policy Rules (continued)

Application/Class Map	WAAS Action	Destination Ports
imap	LZ+TFO+DRE	143
imap3	LZ+TFO+DRE	220
imaps	TFO	993
iso-tsap	LZ+TFO+DRE	102
lotusnote	LZ+TFO+DRE	1352
MAPI[1] (*monitored*)	LZ+TFO+DRE+ MAPI accelerator	UUID:a4f1db00-ca47-1067-b31f-00dd0106 62da
MDaemon	LZ+TFO+DRE	3000, 3001
MS-Exchange-Directory-NSPI[1]	Passthrough	UUID:f5cc5a18-4264-101a-8c59-08002b2f 8426
MS-Exchange-Directory-RFR[1]	Passthrough	UUID:1544f5e0-613c-11d1-93df-00c04fd7 bd09
NNTP (*monitored*)	LZ+TFO+DRE	119
nntps (*monitored*)	TFO	563
openmail	LZ+TFO+DRE	5755, 5757, 5766, 5767, 5768, 5729
pcmail-srv	LZ+TFO+DRE	158
pop3	LZ+TFO+DRE	110
pop3s	LZ+TFO+DRE	995
QMTP	TFO	209
smtp (*monitored*)	LZ+TFO+DRE	25
smtps	TFO	465
Enterprise-Applications *(monitored)*		
MS-GROOVE	TFO	2492
SAP (*monitored*)	LZ+TFO+DRE	3200–3204, 3206–3219, 3221–3224, 3226–3259, 3261–3263, 3265–3267, 3270–3282, 3284–3305, 3307–3351, 3353–3388, 3390–3399, 3600–3658, 3662–3699
Siebel	LZ+TFO+DRE	2320, 2321, 8448
File-System *(monitored)*		
afpovertcp	LZ+TFO+DRE	548
afs3	LZ+TFO+DRE	7000–7009
ncp	LZ+TFO+DRE	524
NFS	LZ+TFO+DRE+ NFS accelerator	2049
sunrpc	Passthrough	111
File-Transfer *(monitored)*		
BFTP	LZ+TFO+DRE	152

Table A-1 Predefined Traffic Policy Rules (continued)

Application/Class Map	WAAS Action	Destination Ports
ftp (*monitored*)	Passthrough	21
ftp-data[2]	LZ+TFO+DRE	20 (source port)
ftps	TFO	990
ftps-data[2]	Passthrough	989 (source port)
sftp	LZ+TFO+DRE	115
TFTP	LZ+TFO+DRE	69
TFTPS	TFO	3713
Instant Messaging		
AOL	Passthrough	5190–5193
Apple-iChat	Passthrough	5297, 5298
ircs	Passthrough	994
ircu	Passthrough	531, 6660–6665, 6667–6669
msnp	Passthrough	1863, 6891–6900
sametime	Passthrough	1533
talk	Passthrough	517
xmpp-client	Passthrough	5222
xmpp-server	Passthrough	5269
Yahoo-Messenger	Passthrough	5000, 5001, 5050, 5100
Name Services		
DNS	Passthrough	53
isns	Passthrough	3205
nameserver	Passthrough	42
netbios	Passthrough	137
svrloc	Passthrough	427
WINS (*monitored*)	Passthrough	1512
Other		
Basic-TCP-services	Passthrough	1–19
BGP	Passthrough	179
corba-iiop-ssl	Passthrough	684
epmap (*monitored*)	TFO, EPM accelerator	135
msmq	LZ+TFO+DRE	1801, 2101, 2103, 2105
NTP	Passthrough	123

Table A-1 Predefined Traffic Policy Rules (continued)

Application/Class Map	WAAS Action	Destination Ports
Other-Secure	Passthrough	261, 448, 695, 994, 2252, 2478, 2479, 2482, 2484, 2679, 2762, 2998, 3077, 3078, 3183, 3191, 3220, 3410, 3424, 3471, 3496, 3509, 3529, 3539, 3660, 3661, 3747, 3864, 3885, 3896, 3897, 3995, 4031, 5007, 7674, 9802, 12109
ssc-agent	LZ+TFO+DRE	2847, 2848, 2967, 2968, 38037, 38292
Unclassified	LZ+TFO+DRE	
P2P *(monitored)*		
BitTorrent	Passthrough	6881–6889, 6969
eDonkey	Passthrough	4661, 4662
Gnutella	Passthrough	5634, 6346–6349, 6355
Grouper	Passthrough	8038
HotLine	Passthrough	5500–5503
Kazaa	Passthrough	1214
Laplink-ShareDirect	Passthrough	2705
Napster	Passthrough	6666, 6677, 6688, 6700, 7777, 8875
Qnext	Passthrough	44, 5555
SoulSeek	Passthrough	2234, 5534
WASTE	Passthrough	1337
WinMX	Passthrough	6699
Printing *(monitored)*		
hp-pdl-datastr	LZ+TFO+DRE	9100
IPP	LZ+TFO+DRE	631
printer	LZ+TFO+DRE	515
print-srv	LZ+TFO+DRE	170
xprint-server	LZ+TFO+DRE	8100
Remote-Desktop *(monitored)*		
Altiris-CarbonCopy	Passthrough	1680
citrixadmin	LZ+TFO+ DRE-unidirectional	2513
citrixima	LZ+TFO+ DRE-unidirectional	2512
citrximaclient (*monitored*)	LZ+TFO+DRE	2598
ControlIT	TFO	799
Danware-NetOp	TFO	6502
ica (*monitored*)	LZ+TFO+DRE	1494

Table A-1 Predefined Traffic Policy Rules (continued)

Application/Class Map	WAAS Action	Destination Ports
laplink	LZ+TFO+ DRE-unidirectional	1547
Laplink-surfup-HTTPS	TFO	1184
ms-wbt-server (*monitored*)	TFO	3389
net-assistant	Passthrough	3283
netrjs-3	TFO	73
pcanywheredata	TFO	5631, 5632, 65301
radmin-port	TFO	4899
Remote-Anything (*monitored*)	TFO	3999, 4000
timbuktu	TFO	407
timbuktu-srv	TFO	1417–1420
Vmware-VMConsole	TFO	902
VNC (*monitored*)	TFO	5800–5809, 5900–5909
x11	TFO	6000–6063
Replication *(monitored)*		
Double-Take	LZ+TFO+ DRE-unidirectional	1100, 1105
EMC-Celerra-Replicator	LZ+TFO+ DRE-adaptive	8888
MS-AD-Replication[1]	LZ+TFO+DRE	UUID:e3514235-4b06-11d1-ab04-00c04fc2 dcd2
ms-content-repl-srv	TFO	507, 560
MS-FRS[1]	LZ+TFO+DRE	UUID:f5cc59b4-4264-101a-8c59-08002b2f 8426
netapp-snapmirror	LZ+TFO+ DRE-adaptive	10565-10569
pcsync-http	LZ+TFO+DRE	8444
pcsync-https	TFO	8443
rrac	TFO	5678
Rsync (*monitored*)	LZ+TFO+ DRE-unidirectional	873
SQL *(monitored)*		
gds_db	LZ+TFO+DRE	3050
IBM-DB2	LZ+TFO+DRE	523
intersys-cache	LZ+TFO+DRE	1972
ms-olap4	TFO	2383
ms-sql-m	LZ+TFO+DRE	1434

Table A-1 Predefined Traffic Policy Rules (continued)

Application/Class Map	WAAS Action	Destination Ports
MS-SQL-RPC[1]	LZ+TFO+DRE	UUID:3f99b900-4d87-101b-99b7-aa000400 7f07
ms-sql-s (*monitored*)	LZ+TFO+DRE	1433
MySQL	LZ+TFO+DRE	3306
Oracle	LZ+TFO+DRE	66
orasrv	LZ+TFO+DRE	1521, 1525
Pervasive-SQL	LZ+TFO+DRE	1583
PostgreSQL	LZ+TFO+DRE	5432
sqlexec	LZ+TFO+DRE	9088, 9089
sql-net	LZ+TFO+DRE	150
sqlserv	LZ+TFO+DRE	118
sqlsrv	LZ+TFO+DRE	156
ssql	LZ+TFO+DRE	3352
sybase-sqlany	LZ+TFO+DRE	1498, 2439, 2638, 3968
UniSQL	LZ+TFO+DRE	1978, 1979
SSH		
SSH (*monitored*)	TFO	22
SSL *(monitored)*		
HTTPS (*monitored*)	TFO	443
Storage *(monitored)*		
EMC-SRDFA-IP	LZ+TFO+DRE	1748
FCIP	LZ+TFO	3225
iFCP	LZ+TFO+DRE	3420
iscsi	LZ+TFO+DRE	3260
Streaming *(monitored)*		
Liquid-Audio	LZ+TFO+ DRE-unidirectional	18888
ms-streaming (*monitored*)	LZ+TFO+ DRE-unidirectional	1755
RTSP (*monitored*)	LZ+TFO+ DRE-unidirectional +Video accelerator	554, 8554
Systems-Management *(monitored)*		
BMC-Patrol	Passthrough	6161, 6162, 6767, 6768, 8160, 8161, 10128
eTrust-policy-Compliance	TFO	1267
flowmonitor	LZ+TFO	7878
HP-OpenView	Passthrough	7426–7431, 7501, 7510

Table A-1 Predefined Traffic Policy Rules (continued)

Application/Class Map	WAAS Action	Destination Ports
LANDesk	LZ+TFO+DRE	9535, 9593–9595
NetIQ	Passthrough	2220, 2735, 10113–10116
Netopia-netOctopus	Passthrough	1917, 1921
netviewdm	Passthrough	729–731
novadigm	LZ+TFO+DRE	3460, 3461, 3464
novell-zen	LZ+TFO+DRE	1761–1763, 2037, 2544, 8039
objcall	LZ+TFO+DRE	94, 627, 1965, 1580, 1581
WBEM	Passthrough	5987–5990
Version-Management *(monitored)*		
Clearcase	LZ+TFO+DRE	371
cvspserver	LZ+TFO+DRE	2401
VPN		
L2TP	TFO	1701
OpenVPN	TFO	1194
PPTP	TFO	1723
Web *(monitored)*		
HTTP (*monitored*)	LZ+TFO+DRE+ HTTP accelerator	80, 3128, 8000, 8080, 8088
soap-http	LZ+TFO+ DRE-adaptive	7627

1. These classifiers use the EPM service in WAAS to accelerate traffic. EPM-based applications do not have predefined ports so the application's UUID must be used to identify the traffic.
2. These classifiers identify the source port instead of the destination port.

APPENDIX **B**

Transaction Log Format

You can use the transaction logging feature to log individual TCP transactions for a WAAS device. For information on configuring transaction logging, see the "Configuring Transaction Logging" section on page 17-53.

TFO transaction logs are kept on the local disk in the directory /local1/logs/tfo.

There are several kinds of transaction log messages that have different templates, as follows

- Optimized Flow Start message:

  Time_Stamp :Conn_ID :Src_IP :Src_Port :Dst_IP :Dst_Port :OT :Log_type :Conn_type :Peer_ID :App_map_name :App_name :App_classifier_name :Flag_directed_mode :TFO_cfgd_policy :TFO_drvd_policy :TFO_peer_policy :TFO_neg_policy :TFO_applied_policy :TFO_reject_reason :AO_cfgd_policy :AO_drvd_policy :AO_neg_policy :AO_reject_reason :SSL_reject_reason :DSCP :Link_rtt

- Optimized Flow End Message:

  Time_Stamp :Conn_ID :Src_IP :Src_Port :Dst_IP :Dst_Port :OT :Log_type :Conn_type :AO_neg_policy :Original_bytes_read :Original_bytes_written :Optimized_bytes_read :Optimized_bytes_written

- Pass Through Flow Message:

  Time_Stamp :Src_IP :Src_Port :Dst_IP :Dst_Port :BP :Bypass_Reason :TFO_cfgd_policy :TFO_drvd_policy :TFO_peer_policy :TFO_reject_reason :AO_cfgd_policy :AO_drvd_policy :AO_reject_reason

- Optimized Flow TFO End Message:

  Time_Stamp :Conn_ID :Src_IP :Src_Port :Dst_IP :Dst_Port :SODRE :END :Original_bytes_read :Original_bytes_written :Optimized_bytes_read :Optimized_bytes_written :Conn_close_state

- System Restart Message:

  Time_Stamp :0 :0 :0 :0 :0 :RESTART

Table B-1 describes the fields found in the transaction log messages.

Table B-1 Transaction Log Field Descriptions

Field	Description
Time_Stamp	Time stamp indicating when the log message was generated.
Conn_ID	A unique identifier for the connection.
Src_IP, Src_Port	Source IP address and port number for the connection.

Table B-1 Transaction Log Field Descriptions (continued)

Field	Description
Dst_IP, Dst_Port	Destination IP address and port number for connection.
OT	Indicates an optimized connection.
BP	Indicates a pass-through connection.
SODRE	Indicates a log message generated by TFO.
Log_type	START or END indicates the start or end of the flow.
Conn_type	Type of connection: INTERNAL CLIENT–locally initiated connection from the WAE, EXTERNAL CLIENT–WAE acting as branch device for the connection, INTERNAL SERVER–locally terminated connection at the WAE, EXTERNAL SERVER–WAE acting as data center device for the connection.
Peer_ID	Device ID of the peer WAE.
App_map_name	Map name.
App_classifier_name	Classifier name.
App_name	Application name.
Flag_directed_mode	T (true) indicates a directed mode connection, F (false) otherwise.
TFO_cfgd_policy	The TFO configured policy on the local device.
TFO_drvd_policy	The TFO derived policy on the local device based on the configured and dynamic conditions. This policy is used to negotiate with the peer WAE.
TFO_peer_policy	The TFO derived policy on the peer that is sent to the local device.
TFO_neg_policy	The TFO negotiated policy, which is the lowest common policy between the derived and peer policies.
TFO_applied_policy	The final policy applied to the connection. After the connection has been established, policy changes may be made to the connection based on the data on the connection, thus the applied policy can differ from the negotiated policy.
TFO_reject_reason	Indicates the reason for a rejected connection. “None” indicates the reject reason is not set.
AO_cfgd_policy	The application accelerator configured on the local device. This is derived from the accelerator configured in the corresponding policy.
AO_drvd_policy	The application accelerator derived policy on the local device.
AO_neg_policy	The application accelerator negotiated policy, which is the lowest common policy between the derived and peer policies.
AO_reject_reason	Indicates the reason an application accelerator rejected the connection. “None” indicates the reject reason is not set.
SSL_reject_reason	Indicates the reason the SSL accelerator rejected the connection. “None” indicates the reject reason is not set.
DSCP	Differentiated Services Code Point value set on the outgoing connection.
Link_rtt	Link round trip time in milliseconds.
Original_bytes_read	Bytes read on the original side of the connection.

Table B-1 Transaction Log Field Descriptions (continued)

Field	Description
Original_bytes_written	Bytes written on the original side of the connection.
Optimized_bytes_read	Bytes read on the optimized side of the connection.
Optimized_bytes_written	Bytes written on the optimized side of the connection.
RESTART	Indicates that the WAE was reloaded and the transaction log process was started.

Here are some examples of transaction log messages:

Fully Optimized on both sides (with SSL rejection)

```
Fri Jan 30 03:15:41 2009 :43 :2.57.223.130 :4808 :2.57.223.2 :443 :OT :START :EXTERNAL CLIENT
:00.14.5e.95.4c.85 :basic :SSL :HTTPS :F :(TFO) (TFO) (TFO) (TFO) (TFO) :<None> :(None) (None) (None) :<None>
:<Keepalive Timeout>  :0 :0
Fri Jan 30 03:15:41 2009 :43 :2.57.223.130 :4808 :2.57.223.2 :443 :SODRE :END :0 :0 :0 :0 :0
Fri Jan 30 03:15:41 2009 :43 :2.57.223.130 :4808 :2.57.223.2 :443 :OT :END :EXTERNAL CLIENT :(None) :284 :806
:806 :28
```

Fully Optimized on both sides

```
Mon Feb  2 14:31:21 2009 :16 :2.75.52.131 :4374 :2.75.52.3 :80 :OT :START :EXTERNAL CLIENT :00.14.5e.83.8c.cf
:basic :Web :HTTP :F :(DRE,LZ,TFO) (DRE,LZ,TFO) (DRE,LZ,TFO) (DRE,LZ,TFO) (DRE,LZ,TFO) :<None> :(HTTP) (HTTP)
(HTTP) :<None> :<None>  :0 :0
Mon Feb  2 14:31:26 2009 :16 :2.75.52.131 :4374 :2.75.52.3 :80 :SODRE :END :370 :173 :299 :429 :0
Mon Feb  2 14:31:26 2009 :16 :2.75.52.131 :4374 :2.75.52.3 :80 :OT :END :EXTERNAL CLIENT :(HTTP) :0 :0 :299
:429
```

Optimized with only DRE enabled

```
Mon Feb  2 14:48:31 2009 :27 :2.75.52.131 :4389 :2.75.52.2 :80 :OT :START :EXTERNAL CLIENT :00.14.5e.83.8c.cf
:basic :Web :HTTP :F :(DRE,TFO) (DRE,TFO) (DRE,LZ,TFO) (DRE,TFO) (DRE,TFO) :<None> :(HTTP) (HTTP) (HTTP)
:<None> :<None>  :0 :0
Mon Feb  2 14:48:36 2009 :27 :2.75.52.131 :4389 :2.75.52.2 :80 :SODRE :END :246 :468 :636 :405 :0
Mon Feb  2 14:48:36 2009 :27 :2.75.52.131 :4389 :2.75.52.2 :80 :OT :END :EXTERNAL CLIENT :(HTTP) :0 :0 :636
:405
```

Optimized with only LZ enabled

```
Mon Feb  2 14:39:12 2009 :20 :2.75.52.131 :4379 :2.75.52.3 :80 :OT :START :EXTERNAL CLIENT :00.14.5e.83.8c.cf
:basic :Web :HTTP :F :(LZ,TFO) (LZ,TFO) (DRE,LZ,TFO) (LZ,TFO) (LZ,TFO) :<None> :(HTTP) (HTTP) (HTTP) :<None>
:<None>  :0 :0
Mon Feb  2 14:39:17 2009 :20 :2.75.52.131 :4379 :2.75.52.3 :80 :SODRE :END :370 :173 :219 :295 :0
Mon Feb  2 14:39:17 2009 :20 :2.75.52.131 :4379 :2.75.52.3 :80 :OT :END :EXTERNAL CLIENT :(HTTP) :0 :0 :219
:295
```

Optimized with both DRE and LZ disabled

```
Mon Feb  2 14:49:36 2009 :28 :2.75.52.131 :4390 :2.75.52.2 :80 :OT :START :EXTERNAL CLIENT :00.14.5e.83.8c.cf
:basic :Web :HTTP :F :(TFO) (TFO) (DRE,LZ,TFO) (TFO) (TFO) :<None> :(HTTP) (HTTP) (HTTP) :<None> :<None>  :0
:0
```

```
Mon Feb  2 14:49:41 2009 :28 :2.75.52.131 :4390 :2.75.52.2 :80 :OT :END :EXTERNAL CLIENT :(HTTP) :0 :0 :468
:246
```

Pass-Through Connection

```
Thu Jul 24 03:09:34 2008 :2.75.52.130 :40027 :2.75.52.2 :80 :BP :GLB_CFG :(DRE,LZ,TFO) (None) (None) :<Global
Config> :(HTTP) (None) :<Global Config>
```

System Restart

```
Sun Oct 25 17:46:32 2009 :0 :0 : 0 :0 :0 :RESTART
```