Cisco Wide Area Application Services Configuration Guide

Software Version 4.2.1
June 22, 2010

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-21609-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

Cisco Wide Area Application Services Configuration Guide
© 2006-2010 Cisco Systems, Inc. All rights reserved.

CONTENTS

Preface xix

PART 1 **WAAS Introduction and Planning**

CHAPTER 1 **Introduction to Cisco WAAS** 1-1

About Cisco WAAS 1-1
Cisco WAAS Overcomes Common WAN Challenges 1-2
Traffic Optimization Process 1-2
Key Services of Cisco WAAS 1-4
TFO Optimization 1-4
Windows Scaling 1-4
TCP Initial Window Size Maximization 1-4
Increased Buffering 1-5
Selective Acknowledgment 1-5
BIC TCP 1-5
Compression 1-5
Application-Specific Acceleration 1-6
File Services for Desktop Applications 1-7
WAAS Print Services 1-8
Virtualization 1-8
Overview of the WAAS Interfaces 1-9
WAAS Central Manager GUI 1-9
Accessing the WAAS Central Manager GUI 1-9
Components of the WAAS Central Manager GUI 1-10
WAAS Central Manager Navigation Pane 1-12
WAAS Central Manager Taskbar Icons 1-12
WAE Device Manager GUI 1-15
WAAS Print Services Administration GUI 1-16
WAAS CLI 1-16
Benefits of Cisco WAAS 1-17
Preservation of Source TCP/IP Information 1-17
Autodiscovery of WAAS Devices 1-17
Centralized Network Monitoring and Management 1-17
Optimized Read and Write Caching 1-18
WCCP Support 1-19

PBR Support 1-20
Inline Interception Support 1-20
Failure Resiliency and Protection 1-20
Namespace Support 1-21
RAID Compatibility 1-21
Streamlined Security 1-21
SNMP Support 1-22

CHAPTER 2 **Planning Your WAAS Network** 2-1

Checklist for Planning Your WAAS Network 2-1
Planning Checklist 2-2
Site and Network Planning 2-4
Windows Network Integration 2-5
Data Center WAE Integration 2-5
Branch WAE Integration 2-6
UNIX Network Integration 2-6
WAFS-Related Ports in a WAAS Environment 2-6
Port 4050 2-7
Ports 139 and 445 2-7
Ports 88 and 464 2-7
Port 50139 2-7
Firewalls and Directed Mode 2-7
Firewalls and Standby Central Managers 2-8
About Autoregistration and WAEs 2-8
Selecting Static IP Addresses or Using Interface-Level DHCP 2-10
Identifying and Resolving Interoperability Issues 2-10
Interoperability and Support 2-10
Unicode Support for the WAAS GUI Interfaces 2-11
Unicode Support Limitations 2-11
WAAS and Cisco IOS Interoperability 2-11
WAAS Support of the Cisco IOS QoS Classification Feature 2-12
WAAS Support of the Cisco IOS NBAR Feature 2-12
WAAS Support of the Cisco IOS Marking 2-14
WAAS Support of the Cisco IOS Queuing 2-14
WAAS Support of the Cisco IOS Congestion Avoidance 2-14
WAAS Support of the Cisco IOS Traffic Policing and Rate Limiting 2-14
WAAS Support of the Cisco IOS Signaling 2-14
WAAS Support of the Cisco IOS Link-Efficiency Operations 2-14
WAAS Support of the Cisco IOS Provisioning, Monitoring, and Management 2-14

WAAS and Management Instrumentation 2-15
WAAS and MPLS 2-15
WAAS Compatibility with other Cisco Appliances and Software 2-15
WAAS Devices and Device Mode 2-16
Calculating the Number of WAAS Devices Needed 2-17
Supported Methods of Traffic Redirection 2-18
Advantages and Disadvantages of Using Inline Interception 2-18
Advantages and Disadvantages of Using WCCP-Based Routing 2-19
Advantages and Disadvantages of Using PBR 2-20
Configuring WCCP or PBR Routing for WAAS Traffic 2-21
Configuring WAEs as Promiscuous TCP Devices in a WAAS Network 2-23
Using Tertiary Interfaces or Subinterfaces to Connect WAEs to Routers 2-24
Access Lists on Routers and WAEs 2-24
IP ACLs on WAEs 2-25
Interception ACLs on WAEs 2-25
Static Bypass Lists on WAEs 2-25
WAAS Login Authentication and Authorization 2-26
WAAS Administrator Accounts 2-26
Logically Grouping Your WAEs 2-27
Data Migration Process 2-28

PART 2 **Installing and Configuring WAAS**

CHAPTER 3 **Using Device Groups and Device Locations** 3-1
About Device and Baseline Groups 3-1
Working with Device Groups 3-2
Creating a Device Group 3-2
Creating a New Device Group 3-3
Configuring the Settings for a Device Group 3-4
Assigning Devices to a Configuration Device Group 3-5
Deleting a Device Group 3-6
Viewing Device Group Assignments 3-6
Viewing the Device Groups List 3-7
Enabling or Disabling Device Group Overlap 3-7
Overriding Group Configuration Settings 3-8
Forcing Device Group Settings on All Devices in the Group 3-8
Selecting Device Group Precedence 3-8
Overriding the Device Group Settings on a Device 3-9

Understanding the Impact of Assigning a Device to Multiple Device Groups 3-10
Working with Baseline Groups 3-10
Configuring the Default Baseline Groups 3-11
Customizing the Baseline Group Settings 3-11
Configuring the Service Settings for a Baseline Group 3-12
Switching the Baseline Group for a Service 3-13
Working with Device Locations 3-14
Creating Locations 3-14
Deleting Locations 3-15
Viewing the Location Tree 3-15

CHAPTER 4 **Configuring Traffic Interception** 4-1
Request Redirection Methods 4-1
Using WCCP to Transparently Redirect TCP Traffic to WAEs 4-3
Guidelines for Configuring WCCP 4-5
Guidelines for File Server Access Methods 4-6
Configuring Advanced WCCP Features on a WCCP-Enabled Router 4-7
Configuring a Router to Support WCCP Service Groups 4-7
Configuring IP Access Lists on a Router 4-10
Setting a Service Group Password on a Router 4-11
Configuring a Loopback Interface on the Router 4-12
Configuring Router QoS for WCCP Control Packets 4-12
Managing WCCP Configurations for WAEs 4-12
Load Balancing and WAEs 4-13
Packet-Forwarding Methods 4-15
Reasons for Packet Rejection and Return 4-16
Layer 3 GRE as a Packet-Forwarding Method 4-17
Layer 2 Redirection as a Packet-Forwarding Method 4-17
WCCP Flow Redirection on WAEs 4-18
Viewing or Modifying the WCCP Settings on WAEs 4-18
Creating a WCCP Service Mask for an Existing WCCP Service 4-23
Viewing or Modifying WCCP Service Masks for WAEs 4-23
Viewing a WCCP Router List Configuration for WAEs 4-24
Modifying the Configuration of WCCP Router Lists for WAEs 4-25
Deleting a WCCP Router List from WAEs 4-25
Defining Additional WCCP Router Lists on WAEs 4-26
Configuring WAEs for a Graceful Shutdown of WCCP 4-27
Configuring Static Bypass Lists for WAEs 4-28
Aggregating Static Bypass Lists 4-29

Configuring an Interception Access Control List 4-29
Configuring Egress Methods for Intercepted Connections 4-30
Configuring a GRE Tunnel Interface on a Router 4-32
Multipoint Tunnel Configuration 4-32
Point-To-Point Tunnel Configuration 4-33
Using Policy-Based Routing to Transparently Redirect All TCP Traffic to WAEs 4-33
Methods of Verifying PBR Next-Hop Availability 4-39
Method 1: Using CDP to Verify Operability of WAEs 4-40
Method 2: Using IP SLAs to Verify WAE Operability Using ICMP Echo Verification (Recommended Method) 4-40
Method 3: Using IP SLAs to Verify WAE Operability Using TCP Connection Attempts 4-41
Using Inline Mode to Transparently Intercept TCP Traffic 4-42
Configuring Inline Interface Settings 4-44
Configuring an IP Address on an Inline Interface 4-47
Configuring VLANs for Inline Support 4-49
Clustering Inline WAEs 4-50

CHAPTER 5 **Configuring Network Settings** 5-1
Configuring Network Interfaces 5-1
Configuring a Standby Interface 5-2
Configuring the Primary Standby Interface 5-3
Configuring Multiple IP Addresses on a Single Interface 5-5
Modifying Gigabit Ethernet Interface Settings 5-5
Configuring Port-Channel Settings 5-7
Configuring Interfaces for DHCP 5-8
Configuring a Load-Balancing Method for Interfaces 5-9
Configuring TCP Settings 5-9
Explicit Congestion Notification 5-11
Congestion Windows 5-11
Retransmit Time Multiplier 5-11
TCP Slow Start 5-12
Path MTU Discovery 5-12
Configuring Static IP Routes 5-13
Aggregating IP Routes 5-13
Configuring CDP Settings 5-14
Configuring the DNS Server 5-14
Configuring Windows Name Services 5-15
Configuring Directed Mode 5-16

CHAPTER 6 **Configuring Administrative Login Authentication, Authorization, and Accounting** 6-1

About Administrative Login Authentication and Authorization 6-1

Default Administrative Login Authentication and Authorization Configuration 6-4

Configuring Administrative Login Authentication and Authorization 6-6

Configuring Login Access Control Settings for WAAS Devices 6-7

Configuring Secure Shell Settings for WAAS Devices 6-8

Disabling and Reenabling the Telnet Service for WAAS Devices 6-10

Configuring Message of the Day Settings for WAAS Devices 6-10

Configuring Exec Timeout Settings for WAAS Devices 6-11

Configuring Line Console Carrier Detection for WAAS Devices 6-12

Configuring Remote Authentication Server Settings for WAAS Devices 6-13

Configuring RADIUS Server Authentication Settings 6-13

About TACACS+ Server Authentication Settings 6-15

Configuring TACACS+ Server Settings 6-16

Configuring Windows Domain Server Authentication Settings 6-17

LDAP Server Signing 6-24

Enabling Administrative Login Authentication and Authorization Schemes for WAAS Devices 6-26

Configuring AAA Command Authorization 6-32

Configuring AAA Accounting for WAAS Devices 6-33

Viewing Audit Trail Logs 6-34

CHAPTER 7 **Creating and Managing Administrator User Accounts and Groups** 7-1

Overview of Administrator User Accounts 7-1

Creating and Managing User Accounts 7-2

Overview for Creating an Account 7-2

Working with Accounts 7-3

Creating a New Account 7-4

Modifying and Deleting User Accounts 7-6

Changing the Password for Your Own Account 7-7

Changing the Password for Another Account 7-8

Viewing User Accounts 7-8

Unlocking User Accounts 7-9

Working with Passwords 7-9

Working with Roles 7-10

Creating a New Role 7-11

Assigning a Role to a User Account 7-13

Modifying and Deleting Roles 7-13

Viewing Role Settings 7-14

Working with Domains 7-14

Creating a New Domain 7-15
Adding an Entity to a Domain 7-15
Assigning a Domain to a User Account 7-16
Modifying and Deleting Domains 7-17
Viewing Domains 7-17
Working with User Groups 7-18
Creating a New User Group 7-18
Assigning Roles to a User Group 7-19
Assigning Domains to a User Group 7-19
Modifying and Deleting a User Group 7-20
Viewing User Groups 7-20

CHAPTER 8 **Creating and Managing IP Access Control Lists for WAAS Devices** 8-1
About IP ACLs for WAAS Devices 8-1
Creating and Managing IP ACLs for WAAS Devices 8-2
List of Extended IP ACL Conditions 8-7

CHAPTER 9 **Configuring Other System Settings** 9-1
Modifying Device Properties 9-1
Managing Software Licenses 9-3
Enabling the Inetd RCP and FTP Services 9-4
Configuring Date and Time Settings 9-5
Configuring NTP Settings 9-5
Configuring Time Zone Settings 9-5
Configuring Secure Store Settings 9-10
Secure Store Overview 9-10
Enabling Secure Store Encryption on the Central Manager 9-12
Enabling Secure Store Encryption on a Standby Central Manager 9-13
Enabling Secure Store Encryption on a WAE Device 9-13
Changing the Secure Store Encryption Key and Password 9-14
Resetting Secure Store Encryption on a Central Manager 9-15
Disabling Secure Store Encryption on a WAE Device 9-16
Modifying the Default System Configuration Properties 9-17
Configuring Web Application Filter 9-19
Enabling Web Application Filter 9-20
Security Verification 9-21
Input Validation 9-21
Sanitization 9-21

Configuring Faster Detection of Offline WAAS Devices 9-22
About Faster Detection of Offline Devices 9-23
Configuring Alarm Overload Detection 9-23
Configuring the E-mail Notification Server 9-24

CHAPTER 10 **Using the WAE Device Manager GUI** 10-1
Launching the WAE Device Manager 10-1
A Quick Tour of the WAE Device Manager 10-2
WAE Management Workflow 10-3
Managing a Cisco WAE 10-4
Control Option 10-4
Starting and Stopping Components 10-5
Registering and Unregistering a WAE 10-6
Backing Up the Configuration Files 10-7
Restoring the Configuration Files 10-7
Configuration Option 10-8
Configuring SNMP Settings 10-8
Viewing Network Settings 10-9
Configuring Windows Authentication 10-10
Defining Notification Settings 10-15
Utilities Option 10-17
Running Support Utilities 10-17
Running the Cache Cleanup Utility 10-18
Running the File Server Rename Utility 10-19
Managing a CIFS Accelerator Device 10-19
Preposition Option 10-20
Terminating a Preposition Task 10-22
Managing a WAFS Core Device 10-22
Managing a WAFS Edge Device 10-23
Monitoring the WAE 10-23
Monitoring Graphs 10-24
Viewing Options 10-24
Monitoring the Cisco WAE Component 10-25
Monitoring a WAFS Core 10-26
Monitoring a Transparent CIFS Accelerator or a WAFS Edge Device 10-28
Viewing WAE Logs 10-32
WAE Logs 10-32
Setting Display Criteria 10-32

Viewing Log Entries 10-33
Saving Log File Information 10-33
Viewing Cisco WAE Logs 10-34

PART 3 Configuring WAAS Services

CHAPTER 11 Configuring Wide Area File Services 11-1
About File Services 11-1
Overview of File Services Features 11-3
Automatic Discovery 11-4
Automatic Discovery with Legacy Mode 11-4
Prepositioning 11-4
Data Coherency 11-5
Data Concurrency 11-6
File-Locking Process 11-6
Microsoft Interoperability 11-7
Windows Shadow Copy for Shared Folders 11-7
Preparing for File Services 11-8
Using File Services on the NME-WAE 11-9
Configuring File Services 11-9
Configuring the Core Cluster 11-11
Configuring the Edge Devices 11-14
Setting Up File Servers to Export to the Edge WAE Cache 11-17
Registering a File Server with the WAAS Central Manager 11-17
Importing File Server Definitions Using a CSV File 11-18
Assigning a Core Cluster to a Registered File Server 11-20
Creating Dynamic Shares 11-21
Creating a Connection Between a Core Cluster and Edge WAEs 11-23
Creating a Preposition Directive 11-25
Creating a New Preposition Directive 11-27
Assigning Edge Devices to a Preposition Directive 11-31
Creating a New Preposition Schedule 11-32
Managing File Services 11-33
Checking the Preposition Status 11-33
Starting and Stopping Preposition Tasks 11-34
Preparing Your WAAS Network for WAN Failures 11-34
About Disconnected Mode 11-35
DNS and Domain Controller Requirements 11-35
Data Availability in Disconnected Mode 11-36

Configuring Disconnected Mode 11-36
Viewing Members of a Core Cluster 11-36
Switching File Services Modes 11-36
Changing from Legacy Mode to Transparent CIFS Accelerator Mode 11-37
Changing from Transparent CIFS Accelerator Mode to Legacy Mode 11-38

CHAPTER 12 **Configuring Application Acceleration** 12-1
About Application Acceleration 12-1
Enabling and Disabling the Global Optimization Features 12-2
Configuring HTTP Acceleration 12-5
About HTTP Metadata Caching 12-7
Configuring MAPI Acceleration 12-8
Configuring Video Acceleration 12-9
Configuring SSL Acceleration 12-11
Preparing to Use SSL Acceleration 12-12
Enabling Secure Store, the Enterprise License, and SSL Acceleration 12-13
Configuring SSL Global Settings 12-14
Configuring a Service Certificate and Private Key 12-15
Working with Cipher Lists 12-19
Working with Certificate Authorities 12-21
Configuring SSL Management Services 12-23
Configuring SSL Peering Service 12-25
Using SSL Accelerated Services 12-26
Creating a New Traffic Application Policy 12-29
Preparing to Create an Application Policy 12-29
Creating an Application Definition 12-30
Creating an Application Policy 12-31
Managing Application Acceleration 12-37
Viewing a List of Applications 12-37
Viewing a Policy Report 12-38
Viewing a Classifier Report 12-38
Restoring Application Policies and Classifiers 12-38
Monitoring Applications 12-39
Defining Default DSCP Marking Values 12-39
Defining the Default DSCP Marking Value 12-39
Modifying the Position of an Application Policy 12-40
Modifying the Acceleration TCP Settings 12-41
Calculating the TCP Buffers for High BDP Links 12-43
Modifying the TCP Adaptive Buffering Settings 12-43

CHAPTER 13 **Configuring and Managing WAAS Legacy Print Services** 13-1

About WAAS Print Services 13-1
Branch Office Printing Topology 13-2
WAAS Print Services 13-3
Print Driver Support and Interoperability 13-3
Printer Clustering 13-4
Print Services Users 13-4
Feature Support 13-4
Planning for Print Services 13-5
Identifying the Print Administration Users 13-5
Obtaining Printer Information 13-6
Planning Worksheets 13-6
Configuring Print Services 13-7
Configuration Checklist 13-7
Preparing your WAE Device and Central Manager for Print Services 13-8
Creating Accounts with Print Admin Privileges 13-10
Enabling Print Services 13-11
Adding a Printer to the WAAS Print Server 13-12
Adding Printer Clusters 13-14
Setting Up the WAAS Central Manager as the Driver Repository 13-16
Installing Print Drivers on Individual WAAS Print Servers 13-18
Distributing Drivers to the WAAS Print Servers 13-20
Distributing a Single Driver to Multiple Devices or Groups 13-20
Distributing Multiple Drivers to a Single Device or Group 13-21
Verifying Print Driver Distribution 13-21
Associating a Driver with your Printer 13-22
Initializing Print Drivers 13-23
Adding the WAAS Print Server to Your Branch Office Clients 13-24
Managing Print Services 13-25
Viewing Print Server Details 13-25
Configuring Aggregate Settings 13-27
Using the Print Services Administration GUI 13-28
Opening the Print Services Administration GUI 13-28
Adding a Printer 13-29
Modifying the Printer Configuration 13-29
Enabling Print Banners 13-31
Setting Up Print Clusters 13-32
Viewing Print Jobs 13-32
Viewing the Print Server Page Log 13-33

Troubleshooting Print Services 13-34
General Known Issues 13-34
Login and Access Problems 13-35
Avoiding Print Problems 13-35
Understanding Interactions Between the WAAS Central Manager and the WAAS CLI 13-36

CHAPTER 14 **Configuring Virtual Blades** 14-1
About Virtual Blades 14-1
Preparing to Use Virtual Blades 14-3
Configuring Virtual Blades 14-4
Installing Paravirtualization Drivers 14-8
Enabling and Disabling Virtual Blades 14-9
Copying a Disk Image to a Virtual Blade 14-10
Backing Up and Restoring a Virtual Blade 14-12

PART 4 **Maintaining, Monitoring, and Troubleshooting your WAAS Network**

CHAPTER 15 **Maintaining Your WAAS System** 15-1
Upgrading the WAAS Software 15-1
Determining the Current Software Version 15-3
Obtaining the Latest Software Version from Cisco.com 15-3
Specifying the Location of the Software File in the WAAS Central Manager GUI 15-4
Upgrading the WAAS Central Manager 15-6
Upgrading Multiple Devices Using Device Groups 15-8
Deleting a Software File 15-9
Backing Up and Restoring your WAAS System 15-9
Backing Up and Restoring the WAAS Central Manager Database 15-10
Backing Up and Restoring a WAE Device 15-11
Using the Cisco WAAS Software Recovery CD-ROM 15-12
Ensuring RAID Pairs Rebuild Successfully 15-16
Recovering the System Software 15-17
Recovering a Lost Administrator Password 15-19
Recovering from Missing Disk-Based Software 15-20
Recovering WAAS Device Registration Information 15-21
Performing Disk Maintenance for RAID-1 Systems 15-22
Replacing Disks in RAID-5 Systems 15-23
Configuring the Central Manager Role 15-24
Converting a WAE to a Standby Central Manager 15-25

Converting a Primary Central Manager to a Standby Central Manager 15-26
Converting a Standby Central Manager to a Primary Central Manager 15-26
Switching Both Central Manager Roles 15-27
Enabling Disk Encryption 15-28
Configuring a Disk Error-Handling Method 15-29
Enabling Extended Object Cache 15-30
Activating All Inactive WAAS Devices 15-31
Rebooting a Device or Device Group 15-32
Performing a Controlled Shutdown 15-33

CHAPTER 16 **Monitoring and Troubleshooting Your WAAS Network** 16-1
Viewing System Information from the System Dashboard Window 16-2
Monitoring Graphs and Charts 16-2
Alarm Panel 16-3
Device Alarms 16-5
Troubleshooting Devices Using Alerts 16-6
Viewing Device Information 16-7
Devices Window 16-7
Device Dashboard Window 16-9
Viewing and Unlocking Device Users 16-10
Customizing a Dashboard or Report 16-10
Adding a Chart 16-12
Configuring Chart Settings 16-12
Chart Descriptions 16-14
Traffic Analysis Charts 16-14
Traffic Summary 16-14
Original Traffic Over Time 16-15
Optimization Charts 16-15
Compression Summary 16-16
Compression Over Time 16-16
Compression by Application Over Time 16-16
Optimized Traffic Over Time 16-17
Traffic Volume and Reduction 16-17
Bandwidth Optimization 16-18
Acceleration Charts 16-18
HTTP 16-18
CIFS 16-21
MAPI 16-25

NFS 16-28
Video 16-31
SSL 16-32
Platform Charts 16-34
Managed Devices Information 16-34
CPU Utilization 16-34
Using Predefined Reports to Monitor WAAS 16-35
Location Level Reports 16-36
Traffic Summary Report 16-36
Optimization Summary Report 16-39
Optimization Details Report 16-39
HTTP Acceleration Report 16-40
Video Acceleration Report 16-41
SSL Acceleration Report 16-42
MAPI Acceleration Report 16-42
NFS Acceleration Report 16-43
Topology Report 16-44
Connections Statistics Report 16-45
CIFS Acceleration Report 16-47
CPU Statistics Report 16-47
Disks Report 16-48
Managing Reports 16-48
Creating Custom Reports 16-49
Viewing and Editing Reports 16-50
Scheduling Reports 16-51
Managing Scheduled Reports 16-52
Configuring Flow Monitoring 16-53
Alarms for Flow Monitoring 16-54
Example Using NetQoS for Flow Monitoring 16-55
Configuring System Logging 16-55
Priority Levels 16-57
Multiple Hosts for System Logging 16-58
Configuring Transaction Logging 16-58
Enabling Transaction Logging 16-58
Transaction Logs 16-60
Viewing the System Message Log 16-61
Viewing the Audit Trail Log 16-62
Viewing the Device Log 16-63
Enabling the Kernel Debugger 16-63

Troubleshooting Using Diagnostic Tests 16-64
Troubleshooting Using the GUI 16-64
Troubleshooting Using the CLI 16-65
Using the show and clear Commands from the WAAS Central Manager GUI 16-65

CHAPTER 17 **Configuring SNMP Monitoring** 17-1
About SNMP 17-1
SNMP Communication Process 17-2
Supported SNMP Versions 17-3
SNMP Security Models and Security Levels 17-3
Supported MIBs 17-4
Downloading MIB Files 17-8
Enabling the SNMP Agent on a WAAS Device 17-8
Checklist for Configuring SNMP 17-8
Preparing for SNMP Monitoring 17-9
Enabling SNMP Traps 17-10
Defining SNMP Traps 17-12
Aggregating SNMP Triggers 17-14
Specifying the SNMP Host 17-14
Specifying the SNMP Community String 17-15
Creating SNMP Views 17-16
Creating an SNMP Group 17-17
Creating an SNMP User 17-18
Configuring SNMP Asset Tag Settings 17-20
Configuring SNMP Contact Settings 17-20
Configuring SNMP Trap Source Settings 17-21

APPENDIX A **Predefined Application Policies** A-1

APPENDIX B **Transaction Log Format** B-1

INDEX

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, page xix
- Document Organization, page xix
- Document Conventions, page xxi
- Related Documentation, page xxii
- Obtaining Documentation and Submitting a Service Request, page xxii

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 Traffic Interception	Describes the WAAS software support for intercepting all TCP traffic in an IP-based network.
Chapter 5	Configuring Network Settings	Describes how to configure basic network settings like DNS and CDP.
Chapter 6	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 7	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 8	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 9	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 10	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 11	Configuring Wide Area File Services	Describes how to configure Wide Area File Services (WAFS), which allows branch office users to more efficiently access data stored at centralized data centers. The WAFS feature overcomes the WAN latency and bandwidth limitations by caching data on WAEs near branch office users.
Chapter 12	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 13	Configuring and Managing WAAS Legacy Print Services	Describes how to configure and manage the WAAS legacy print services feature that allows WAEs to function as print servers in your branch offices.
Chapter 14	Configuring Virtual Blades	Describes how to configure virtual blades, which emulate another computer in your WAAS device.
Chapter 15	Maintaining Your WAAS System	Describes the tasks you may need to perform to maintain your WAAS system.
Chapter 16	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 17	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 Application Policies	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 Installation and Configuration Guide for Windows on a Virtual Blade*
- *Cisco WAAS Troubleshooting Guide for Release 4.1.3 and Later*
- *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*
- *Cisco Wide Area Application Services Online Help*
- *Using the Print Utilities to Troubleshoot and Fix Samba Driver Installation Problems*
- *Regulatory Compliance and Safety Information for the Cisco Wide Area Virtualization Engines*
- *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 512 and 612 Hardware Installation Guide*
- *Cisco Wide Area Application Engine 7326 Hardware Installation Guide*
- *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.

PART 1

WAAS Introduction and Planning

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 appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

This chapter contains the following sections:

- About Cisco WAAS, page 1-1
- Key Services of Cisco WAAS, page 1-4
- Overview of the WAAS Interfaces, page 1-9
- Benefits of Cisco WAAS, page 1-17

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 application policies to determine whether to optimize the traffic or allow it to pass through your network unoptimized.

You use the WAAS Central Manager GUI to centrally configure and monitor the WAEs and application policies in your network. You can also use the WAAS Central Manager GUI to create new application policies 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)

This section contains the following topics:

- Cisco WAAS Overcomes Common WAN Challenges, page 1-2
- Traffic Optimization Process, page 1-2

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.
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

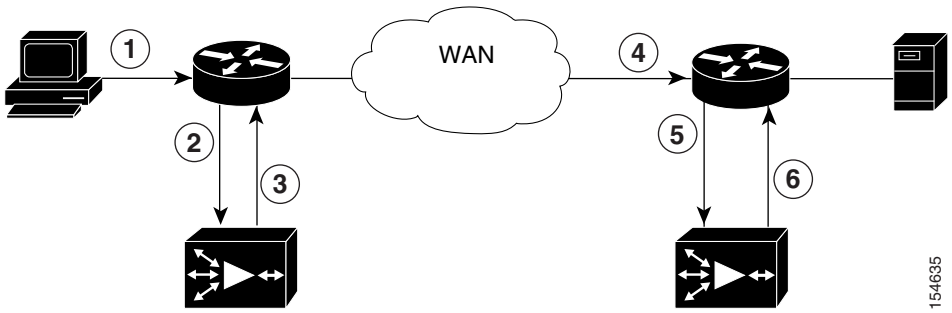

The following steps describe how your WAAS network optimizes a connection between a branch office client and a destination server:

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 a WAE with a Cisco WAE Inline Network Adapter, WAAS can intercept the request directly using inline mode. For more information on inline mode, see the "Using Inline Mode to Transparently Intercept TCP Traffic" section on page 4-42.
3. The branch WAE performs the following actions:

- Examines the parameters in the traffic's TCP headers and then refers to the application policies to determine if the intercepted traffic should be optimized. Information in the TCP header, such as the source and destination IP address, allows the branch WAE to match the traffic to an application policy. For a list of predefined policies, see Appendix A, "Predefined Application Policies."
- 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.

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 application policy 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 1-5
- Application-Specific Acceleration, page 1-6
- File Services for Desktop Applications, page 1-7
- WAAS Print Services, page 1-8
- Virtualization, page 1-8

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 1-4
- TCP Initial Window Size Maximization, page 1-4
- Increased Buffering, page 1-5
- Selective Acknowledgment, page 1-5
- BIC TCP, page 1-5

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.
- WAFS 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.

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 application policies to match specific types of traffic to an application and to determine whether that application traffic should be optimized and accelerated.

The following WAAS application accelerators are available:

- CIFS—Accelerates CIFS traffic exchanged with a remote file server. For more information, see the "File Services for Desktop Applications" section on page 1-7.
- NFS—Accelerates Network File System (NFS) version 3 traffic exchanged with a remote file server. Secure NFS traffic is not accelerated.
- HTTP—Accelerates HTTP 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 2000–2007 clients are supported. Secure connections that use message authentication (signing) or encryption are not accelerated and 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. Server Message Block (SMB) signed traffic is not accelerated. The Windows print accelerator supports Windows 2000 and Windows Server 2003 print servers. It supports clients running Windows 2000, Windows XP, and Windows Vista.

To enable or disable application accelerators, see the "Enabling and Disabling the Global Optimization Features" section on page 12-2.

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 (CIFS accelerator) feature allows a WAE to store remote file server data in its local cache so that the WAE can quickly fulfill a client's data request instead of sending that request over the WAN to the file server. By fulfilling the client's request 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 11, "Configuring Wide Area File Services."

WAAS version 4.1.1 and later supports two mutually exclusive file services modes: a transparent CIFS accelerator mode and a legacy mode. The transparent CIFS accelerator mode requires no core, edge, or connectivity configuration. You configure the legacy mode the same as in WAAS version 4.0.x. These two modes are mutually exclusive. We recommend using the transparent CIFS accelerator mode if you do not need to interoperate with WAAS 4.0.x devices.

Note Legacy mode WAFS is deprecated in WAAS version 4.2.1. It still operates but is not supported and will be removed in a future version. Legacy WAFS users should migrate to the transparent CIFS accelerator.

File Services includes the following features:

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

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 the CIFS accelerator and Windows print acceleration be enabled. For more information, see the "Enabling and Disabling the Global Optimization Features" section on page 12-2.
- 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."
- Legacy WAAS print services—This legacy service also allows you to use the branch WAE as a local print server by using Samba and Common Unix Printing System (CUPS) software that is built into WAAS. For more information, see Chapter 13, "Configuring and Managing WAAS Legacy Print Services."

Note The legacy print services feature is deprecated in WAAS version 4.2.1. It still operates but is not supported and will be removed in a future version. Legacy print services users should migrate to the Windows print accelerator.

All three of 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 1-9
- WAE Device Manager GUI, page 1-15
- WAAS Print Services Administration GUI, page 1-16
- WAAS CLI, page 1-16

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 or device group.
- Create and edit application policies that determine the action that a WAAS device performs when it intercepts specific types of traffic.
- Configure file services and set up file preposition policies.
- 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 file services, print services, or application acceleration 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, page 1-9
- Components of the WAAS Central Manager GUI, page 1-10
- WAAS Central Manager Navigation Pane, page 1-12
- WAAS Central Manager Taskbar Icons, page 1-12

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 7, "Creating and Managing Administrator User Accounts and Groups."

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

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 9-17.

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 after 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 Internet Explorer before restarting 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

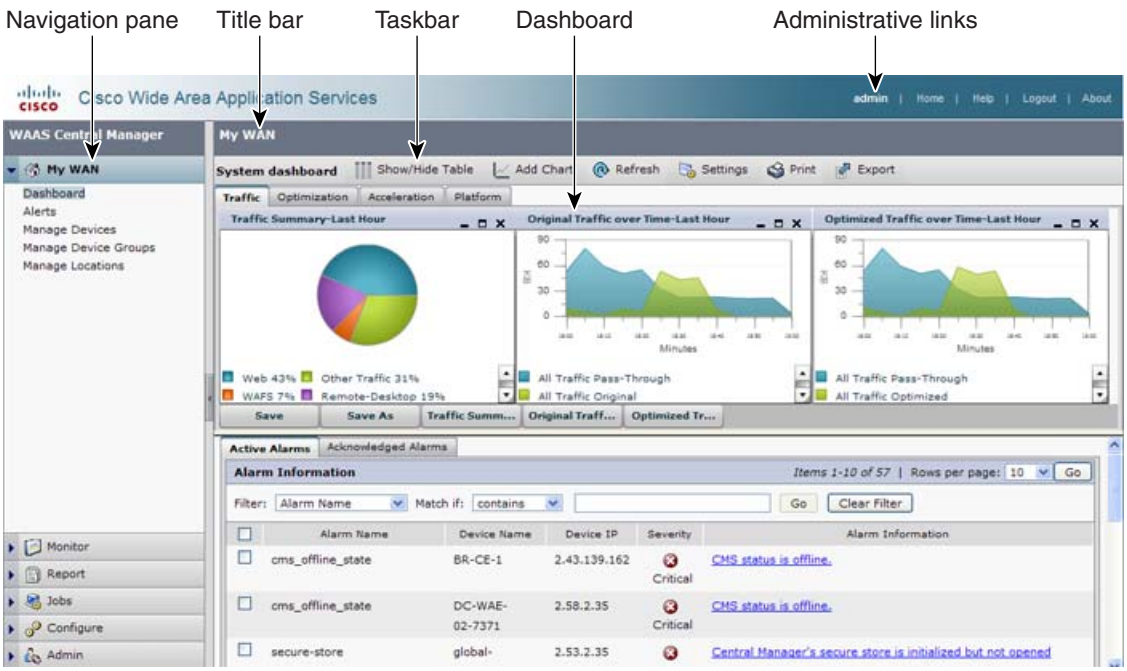

The WAAS Central Manager GUI includes the following main components:

- Navigation pane—Contains drawers that group the WAAS Central Manager functions. For more information, see the "WAAS Central Manager Navigation Pane" section on page 1-12.
- Title bar—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.)

- Taskbar—Contains 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 1-12.
- Dashboard—Displays the main content, which changes depending on the function that is chosen in the navigation pane.
- Administrative links—Includes these navigation links:
  - Home—Displays the System Dashboard (shown in Figure 1-2).
  - Help—Opens a separate window with the WAAS context sensitive help.
  - Logout—Logs out the current user from the WAAS Central Manager.
  - About—Displays the WAAS About screen that shows the version number.

There are two main contexts in the WAAS Central Manager GUI:

- Global Context—You are in this context initially, when no particular device or device group is chosen. The top drawer in the navigation pane reads My WAN, as shown in Figure 1-2.
- Device Context—You are in this context after you have chosen a device or device group in the Manage Devices or Manage Device Groups pages. The top drawer in the navigation pane displays the name of the chosen device or device group, as shown in Figure 1-3.

Figure 1-3 WAAS Central Manager GUI Device Context

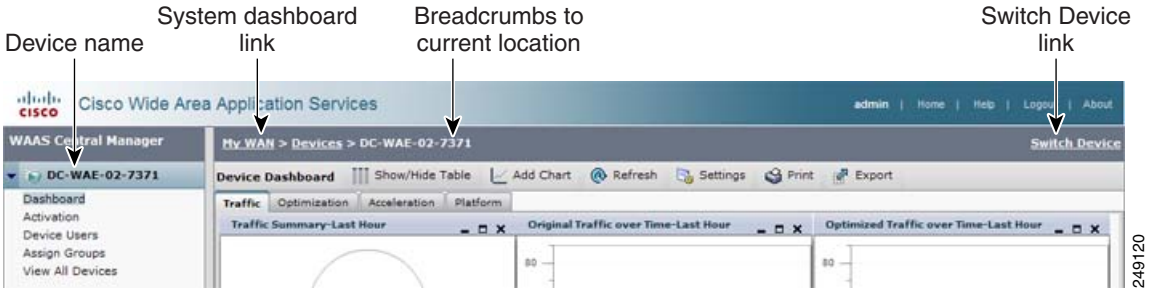

In the device context, the WAAS Central Manager GUI includes the following new items:

- Device name—The top drawer in the navigation pane shows the name of the chosen device or device group.
- System dashboard link—My WAN displays the system dashboard and returns to the global context.
- Breadcrumbs to current location—Displays the path to your current location in the GUI. You can click the Devices link to return to the Manage Devices page. If you are managing a device group, this link is named Device Groups and it returns you to the Manage Device Groups page.
- Switch Device link—Allows you to switch devices and stay on the same functional device context page in the GUI. Clicking this link opens the Switch Devices page, as shown in Figure 1-4. To switch to a different device, select the radio button next to the device and click the **Switch** button.

  If you are managing a device group, the name of this link is Switch DeviceGroup, and it allows you to switch to a different device group, just like for devices.

Figure 1-4 Switch Devices Page

WAAS Central Manager Navigation Pane

The WAAS Central Manager GUI navigation pane contains drawers that serve to group WAAS Central Manager functions. Table 1-2 describes the drawers in the navigation pane.

Drawers may contain different functions when a particular device or device group is selected than when you are in the global context.

Table 1-2 Drawer Descriptions

Drawer	Description
My WAN or *Device name*	Allows you to go to the Dashboard and Alerts displays for your entire WAAS network and allows you to choose a specific device or device group for which to configure WAAS services and general settings. You can also view detailed device information and messages, and manage locations. When you have chosen a specific device or device group, this drawer is named with the device or group name and allows you to activate devices or assign groups or devices. Choose **View All Devices** to go back to the view of all devices.
Monitor	Allows you to see network traffic and other charts and reports to monitor the health and performance of your WAAS network. (This drawer appears only at the global level or when an individual device is chosen, not for a device group.)
Report	Allows you to manage and schedule reports for your WAAS network. (This drawer appears only at the global level when no device or device group is chosen.)
Troubleshoot	Allows you to use troubleshooting tools. (This drawer appears only when a device or device group is chosen, not at the global level.)
Jobs	Allows you to manage software update jobs.
Configure	Allows you to configure the main WAAS services (file, print, and application acceleration) and other settings.
Admin	Allows you to manage user accounts, passwords, licenses, and virtual blades, and view the system log.

WAAS Central Manager Taskbar Icons

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

Table 1-3 Taskbar Icon Descriptions

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, device group, print driver, or file service policy.
(Create)	Creates a new WAAS element, such as a file service policy or an acceleration policy.
(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.
(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.
Devices and Device Group Icons	
(Activate All Inactive WAEs)	Activates all the inactive WAEs in your WAAS network. For more information, see the "Activating All Inactive WAAS Devices" section on page 15-31.
(Force Update)	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. You can view device CLI errors in the System Message window described in the "Viewing the System Message Log" section on page 16-61. The Force Full Database Update icon appears on the Device Dashboard window, described in the "Device Dashboard Window" section on page 16-9.
(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 15-32.

Table 1-3 Taskbar Icon Descriptions (continued)

Taskbar Icon	Function
(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 3-8.
(Apply Defaults)	Applies the default settings to the fields on the window.
(Export Table)	Exports table information into a CSV file.
(Switch Baseline Group)	Allows you to select another device group to associate with the baseline group. For more information, see the "Switching the Baseline Group for a Service" section on page 3-13.
(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 3-9.
(Deactivate Device)	Deactivates a WAE.
(Update Application Statistics)	Updates the application statistics.
(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 in the navigation pane.
(Configure Dashboard Display)	Allows you choose which charts to display in the Device Dashboard window.
Print Services Icons	
(Retry Downloading Failed Drivers)	Attempts to download print drivers that previously failed to be distributed to the WAAS print server or device group. For more information, see Chapter 13, "Configuring and Managing WAAS Legacy Print Services."
(Print Services Administration GUI)	Opens the Print Services Administration GUI for the WAAS print server. For more information about the tasks you can perform from this GUI, see the "Using the Print Services Administration GUI" section on page 13-28.
Acceleration Icons	
(Restore Default Policies and Classifiers)	Restores the default predefined application policies on the device or device group. For more information, see the "Restoring Application Policies and Classifiers" section on page 12-38.
(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 16-44.
(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 12-37.

Table 1-3 Taskbar Icon Descriptions (continued)

Taskbar Icon	Function
System Message Log Icons	
(Truncate Table)	For more information, see the "Viewing the System Message Log" section on page 16-61.

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, the following tasks can only be performed from the WAE Device Manager GUI and not from the WAAS Central Manager GUI:

- Enabling print services on a WAE
- Shutting down device services

For more information about the tasks you can perform from the WAE Manager, see Chapter 10, "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 Example of the WAE Device Manager Window

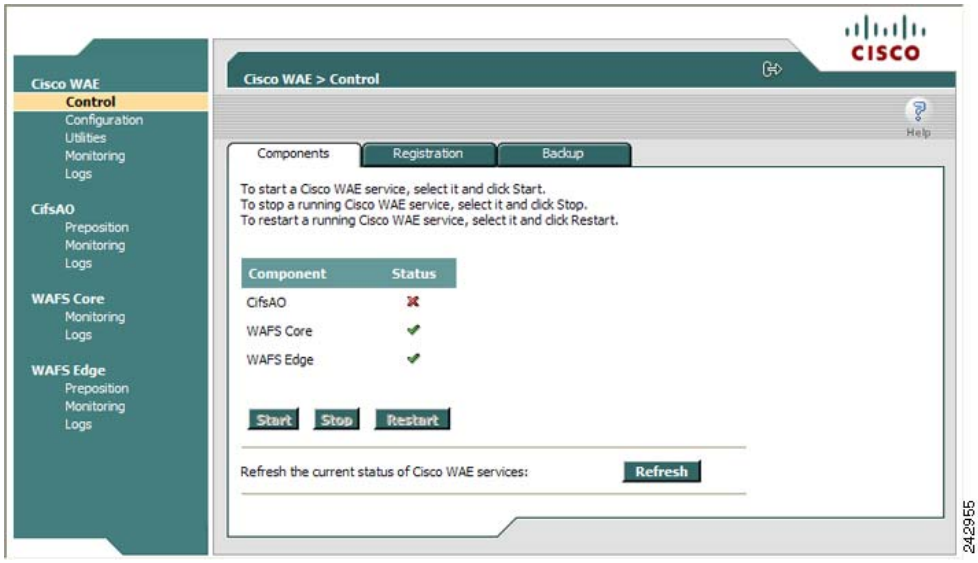

WAAS Print Services Administration GUI

The Print Services Administration GUI is a Web-based interface that allows you to configure an individual WAAS print server and view a list of active and completed print jobs.

You can perform the following common tasks from the Print Services Administration GUI:

- Add a printer to WAAS print server
- Modify the configuration of an existing printer
- Set up print clusters
- View print jobs

You can access the Print Services Administration GUI from the WAAS Central Manager GUI or from the WAE Manager GUI. For more information, see Chapter 13, "Configuring and Managing WAAS Legacy Print Services."

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, page 1-17
- Autodiscovery of WAAS Devices, page 1-17
- Centralized Network Monitoring and Management, page 1-17
- Optimized Read and Write Caching, page 1-18
- WCCP Support, page 1-19
- PBR Support, page 1-20
- Inline Interception Support, page 1-20
- Failure Resiliency and Protection, page 1-20
- Namespace Support, page 1-21
- RAID Compatibility, page 1-21
- Streamlined Security, page 1-21
- SNMP Support, page 1-22

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. If communication between WAEs is disrupted, the software automatically switches into Disconnected Mode, preventing operations that could jeopardize the coherency of files in the network.

Optimized Read and Write Caching

The wide area file services (WAFS) 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. WAFS 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.

In WAFS legacy mode, Cisco WAAS uses its own proprietary adaptation protocol layer over the WAN between the branch WAE and data center WAE, while retaining the standard CIFS protocol at the client and server ends. This proprietary network protocol provides reliable and efficient communication over WANs, especially under high-latency, low-bandwidth conditions.

The Cisco WAAS protocol offers the following benefits:

- Reliability—Maintains its own internal message queuing and ordering, enabling it to overcome transient disconnects, network jitters, and message loss. The Cisco WAAS transport layer handles temporary network failures by reestablishing the connection, then retransmitting requests that did not receive a response on the disconnected socket.
- Efficiency—Supports compound requests, grouping multiple, dependent requests and responses into a single message. The processing of individual calls within a compound message is serialized, enabling the output of one command to be used as input for the next.
- Link utilization optimization—Uses multiple concurrent TCP connections for each branch WAE-to-data center WAE link. Requests and responses may be delivered across any open connection. For example, multiple requests (and responses) for data delivery can be split across multiple connections to increase the effective use of the network in cases of high-latency or high-loss WAN connections, where TCP performance degrades.
- Command prioritization—Assigns high priority to requests from active clients, minimizing the WAN latency experienced by users. Batch tasks (such as preposition, for example) are assigned a lower priority and are performed in the background.
- Bandwidth conservation—Compresses all requests and responses. Before compression, the message is encoded, allowing efficient delivery of both textual and binary data. The protocol layer applies the compression automatically, regardless of the message content.
- Firewall-friendly—Is layered over TCP/IP and uses TCP port 4050. You should configure firewalls to open TCP port 4050 to traffic.

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 4, "Configuring Traffic Interception."

Inline Interception Support

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

The Cisco WAE Inline Network Adapter 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.

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

You can serially cluster WAE devices with the Cisco WAE Inline Network Adapter installed 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.

For more information about inline mode, see the "Using Inline Mode to Transparently Intercept TCP Traffic" section on page 4-42.

Failure Resiliency and Protection

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

If a WAE configured for WAFS 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.

When using legacy mode and communication is interrupted between the Edge WAE and the Core Cluster or between the Core Cluster and the file server, the Cisco WAAS network switches to working in a disconnected state until full communication is restored. If the interruption is brief, the network enters a transient disconnect state, enabling a select number of services and commands for a limited time (typically lasting about one minute), such as read commands for files that are already open.

If the network outage is prolonged, Cisco WAAS switches to a full disconnect state where no services are provided to clients. In this mode, the system denies access to any file (including cached files) until reconnection occurs. From a user viewpoint, the Edge WAE responds as if the network to which it is connected is disconnected.

This approach is required to maintain the security of the data. If a no-service state was not enforced, users connected locally to the file servers can continue working on files, which creates conflicts with other users who may have been working on those files remotely when the network interruption occurred. Cisco WAAS is designed to prevent scenarios that could compromise data coherency and concurrency.

Note Disconnected operation functions only with CIFS legacy mode, not with the CIFS accelerator.

Namespace Support

For legacy CIFS users, there are several ways to access the file servers cached by the Edge WAEs and integrate them within the organizational namespace. One method is to use a prefix, suffix, or alias for a specific site, which creates a unique name for each file server. (Using an alias enables the old name to be retained after replacing the local file server with the new server in the data center.) Another method is to integrate the cached file servers within the DFS namespace as DFS links. When using DFS with legacy CIFS, the DFS site name must be configured manually for each Edge WAE (or edge device group). This information enables DFS to direct user requests correctly. Remote users are directed to file servers through the appropriate Edge WAE, while local users continue to access files directly, without making use of the Edge WAE cache.

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 15, “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.

Cisco WAAS exports parameters based on the following private, read-only MIBs:

- ACTONA-ACTASTOR-MIB.my
- CISCO-CONTENT-ENGINE-MIB

In addition, Cisco WAAS supports the full functionality of each of these standard MIBs, including the setting of traps. 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.

- MIB-2 General Network Statistics (RFC 1213 and 1157)—Contains essential parameters for the basic management of TCP/IP-based networks.
- Host Resources (RFC 1514)
- SNMPv3 MIBs (RFC 2571 through 2576)
- DISMAN-EVENT-MIB (RFC 2981)
- ENTITY-MIB (RFC 2037)

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 Support, see Chapter 17, “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 appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

This chapter contains the following sections:

- Checklist for Planning Your WAAS Network, page 2-1
- Site and Network Planning, page 2-4
- About Autoregistration and WAEs, page 2-8
- Identifying and Resolving Interoperability Issues, page 2-10
- WAAS Devices and Device Mode, page 2-16
- Calculating the Number of WAAS Devices Needed, page 2-17
- Supported Methods of Traffic Redirection, page 2-18
- Access Lists on Routers and WAEs, page 2-24
- WAAS Login Authentication and Authorization, page 2-26
- Logically Grouping Your WAEs, page 2-27
- Data Migration Process, page 2-28

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 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 2-17.
- 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 6, "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 7, "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 6-7.

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 example, Wide Area File Services [WAFS] services, WAAS print services, and WAAS application acceleration). 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 9-3.
- Determine which traffic interception methods to use in your WAAS network (for example, inline mode; WCCP Version 2 or policy-based routing (PBR) for promiscuous mode; NetBIOS for WAFS-only traffic). For more information, see the "Supported Methods of Traffic Redirection" section on page 2-18.

Note WCCP works only with IPv4 networks.

- If you plan to use the 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 2-24.

- 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 2-24.

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 4-39.
- 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 Application Policies."

- Determine the print services configuration. For more information, see Chapter 13, "Configuring and Managing WAAS Legacy Print Services."
- 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.
- Identify the servers, the WAFS file servers that will be used as the target WAFS file servers, and the desired feature set (for example, disconnected mode and home directories).

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 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 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. If you plan to deploy the Wide Area Files Services (WAFs) in legacy mode or enable directed mode, then you must configure the firewall to open port 4050 (UDP for directed mode). However, if you plan to deploy the transparent CIFS accelerator or only generic TCP optimizations and not use directed mode, you do not need to configure the firewall to open port 4050.

This section contains the following topics:

- Windows Network Integration, page 2-5
- UNIX Network Integration, page 2-6
- WAFS-Related Ports in a WAAS Environment, page 2-6
- Firewalls and Directed Mode, page 2-7
- Firewalls and Standby Central Managers, page 2-8

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, page 2-5
- Branch WAE Integration, page 2-6

Note If the integration of WAFs is nontransparent, a WAAS device does not assume Windows server roles on its network, nor does it act as a Domain Controller or master browser in a Windows environment. Another Windows machine should fill these roles in the WAE network. This caveat is not relevant for WCCP or PBR environments because in this situation transparent integration is used.

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).

WAFS-Related Ports in a WAAS Environment

This section describes the Wide Area File Services (WAFS)-related ports used between your clients, WAEs that are functioning as file engines, and CIFS file servers. Most WAFS 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.

Port 4050

Communication between the Core WAE gateway and Edge WAE cache is done over TCP/IP port 4050 by the legacy WAFS feature. If you enable the transparent CIFS accelerator, port 4050 is not used for CIFS.

Ports 139 and 445

If you have only deployed WAFS 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 WAFS 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 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.

Port 50139

If you set up WAAS print services, the print server runs on port 50139. For more information about configuring WAAS print services, see Chapter 13, “Configuring and Managing WAAS Legacy Print Services.”

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 5-16.

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.

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. For WAAS applicances, the first interface is GigabitEthernet 1/0. 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 2-10.)

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 status 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, page 2-10
- WAAS and Cisco IOS Interoperability, page 2-11
- WAAS Compatibility with other Cisco Appliances and Software, page 2-15

Interoperability and Support

Table 2-1 lists the hardware, client, and web browser support for the WAAS software.

Table 2-1 Hardware, Client, Web Browser Support

Hardware support	WAE-512, WAE-612, WAE-674, WAE-7326, WAE-7341, WAE-7371, WAVE-274, WAVE-474, and WAVE-574 appliances, or an NME-WAE, SM-SRE-700, or SM-SRE-900 network module that is installed in specific Cisco routers. You must deploy the WAAS Central Manager on a dedicated device.
Client support	The WAAS software running on a branch WAE interoperates with these CIFS clients: Windows 98/NT 4.0/2000/XP/7 and Windows Server 2003/2008 R2.
Web browser support	The WAAS GUIs require Internet Explorer 5.5 or later to run.

This section contains the following topics:

- Unicode Support for the WAAS GUI Interfaces
- Unicode Support Limitations

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:

- WAAS Support of the Cisco IOS QoS Classification Feature, page 2-12
- WAAS Support of the Cisco IOS NBAR Feature, page 2-12
- WAAS Support of the Cisco IOS Marking, page 2-14

- WAAS Support of the Cisco IOS Queuing, page 2-14
- WAAS Support of the Cisco IOS Congestion Avoidance, page 2-14
- WAAS Support of the Cisco IOS Traffic Policing and Rate Limiting, page 2-14
- WAAS Support of the Cisco IOS Signaling, page 2-14
- WAAS Support of the Cisco IOS Link-Efficiency Operations, page 2-14
- WAAS Support of the Cisco IOS Provisioning, Monitoring, and Management, page 2-14
- WAAS and Management Instrumentation, page 2-15
- WAAS and MPLS, page 2-15

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 2-12, 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 degrade 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 2-12, 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 2-24.
- 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 5-16.

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 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 device needs to use.
- WAAS application accelerator mode—Mode for a WAAS Accelerator, that is data center WAEs and branch WAEs that are running WAAS software.

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.

```
waas-cm(config)# device mode ?
  application-accelerator  Configure device to function as a WAAS Engine.
  central-manager          Configure device to function as a WAAS Central Manager.
```

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)#
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] y
Shutting down all services, will Reload requested by CLI@ttyS0.
Restarting system.
```

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 will not be available.

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 flow 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 Version 2, providing failover capabilities. WCCP Version 2 is the recommended method for larger user populations.
- Number of data center WAEs—Each organization must have at least one data center WAE.

When determining the number of each 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.

  To prevent data loss due to system limitations, WAAS supports a Core WAE cluster for legacy WAFS functionality. This defined group of Core WAEs is used to do the following:

  - Increase the scalability of the capacity of the system.
  - Provide redundancy.

  When using the transparent CIFS accelerator for WAFS, a Core WAE cluster is not needed because the WAAS system automatically routes requests through other WAEs in the event of an outage.

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 installed can operate in inline mode and receive and optimize traffic directly before it passes through the router.

This section contains the following topics:

- Advantages and Disadvantages of Using Inline Interception, page 2-18
- Advantages and Disadvantages of Using WCCP-Based Routing, page 2-19
- Advantages and Disadvantages of Using PBR, page 2-20
- Configuring WCCP or PBR Routing for WAAS Traffic, page 2-21

For detailed information about how to configure traffic interception for your WAAS network, see Chapter 4, “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 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 Cisco WAE Inline Network Adapter contains one or two pairs of LAN/WAN Ethernet ports. If the adapter has two pairs of ports, it can connect to two routers if the network topology requires it. Additionally, some WAE models can have two Cisco WAE Inline Network Adapters installed.

The Cisco WAE Inline Network Adapter 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.

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

You can serially cluster WAE devices with the Cisco WAE Inline Network Adapter installed 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 using hardware accelerated WCCP interception or load balancing with the Cisco Application Control Engine (ACE).

For more information on inline interception, see the "Using Inline Mode to Transparently Intercept TCP Traffic" section on page 4-42.

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

- Dual inline Cisco WAE Inline Network Adapters supported in particular WAE models, providing a total of 8 ports in 4 inline groups.
- Serial inline clustering of two WAEs 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 an Interception Access Control List" section on page 4-29.

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 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 attractive features built-in, for example, automatic failover and load balancing between multiple devices. The WCCP-enabled router monitors the liveness of each WAE 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 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. In the configuration example provided in the *Cisco Wide Area Application Services Quick Configuration Guide,* the CLI is used 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 *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 "Managing WCCP Configurations for WAEs" section on page 4-12. 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 a WCCP-Enabled Router" section on page 4-7.

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 to Transparently Redirect All TCP Traffic to WAEs" section on page 4-33.

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

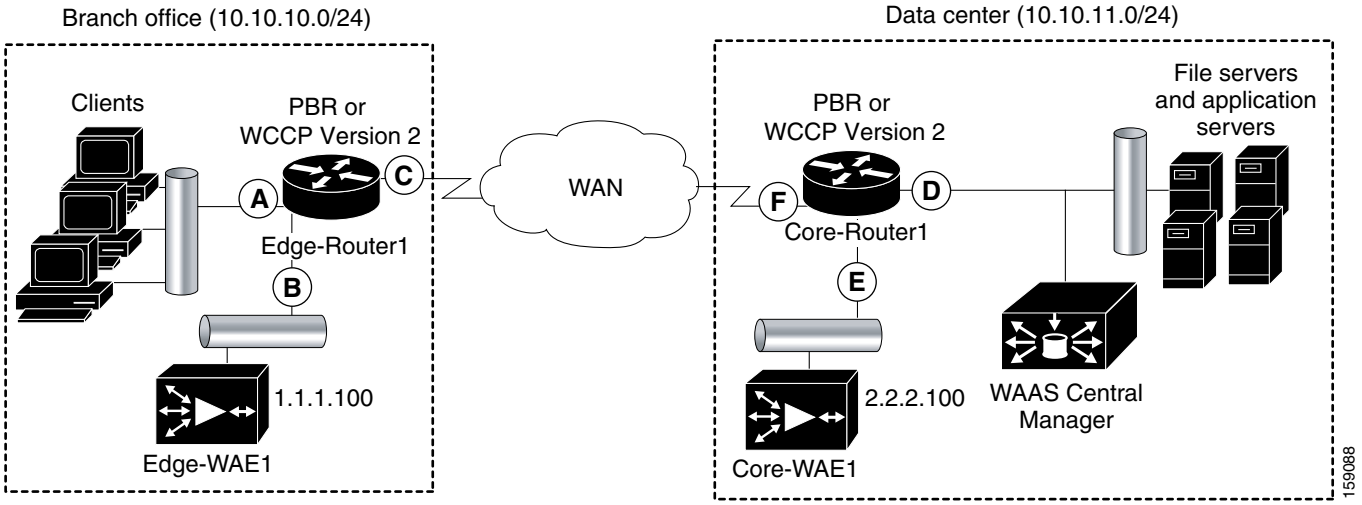

Table 2-2 provides a summary of the router interfaces that you must configure to use PBR or WCCP Version 2 to transparently redirect traffic to a WAE.

Table 2-2 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).

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 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.

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 2-23.

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, page 2-23
- Using Tertiary Interfaces or Subinterfaces to Connect WAEs to Routers, page 2-24

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. The WCCP services 61 and 62 are represented by the canonical name tcp-promiscuous on the WAE, as shown in the following sample output of the WAAS CLI on a branch WAE:

```
Edge-WAE1(config)# wccp ?
  access-list      Configure an IP access-list for inbound WCCP encapsulated
                   traffic
  flow-redirect    Redirect moved flows
  router-list      Router List for use in WCCP services
  shutdown         Wccp Shutdown parameters
  slow-start       accept load in slow-start mode
  tcp-promiscuous  TCP promiscuous mode service
  version          WCCP Version Number
```

WCCP services 61 and 62 are represented by the name TCP Promiscuous in the WAAS Central Manager GUI. (See Figure 4-3.)

Although you can also use the WAAS Central Manager GUI to configure the TCP promiscuous mode service on an individual WAE, we recommend that you use the WAAS CLI to complete the initial basic configuration of the WAE and then use the WAAS Central Manager GUI to make any subsequent configuration changes. By using the WAAS Central Manager GUI to make subsequent configuration changes, you can also apply those changes to groups of WAEs (device groups). 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 or group of WAEs, see the "Managing WCCP Configurations for WAEs" section on page 4-12.

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 Intercepted Connections" section on page 4-30.

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 4-10.

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, page 2-25
- Interception ACLs on WAEs, page 2-25
- Static Bypass Lists on WAEs, page 2-25

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 8, "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 an Interception Access Control List" section on page 4-29.

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 on WAEs

In addition to defining interception ACLs, you can also configure static bypass lists on the WAEs. When you use static bypass, the WAE can bypass the handling of traffic flows between a configurable set of clients and servers. By configuring static bypass entries on the branch WAE, you can control traffic interception without modifying the router configuration. You can also configure access lists on the router to bypass traffic without first redirecting it to the branch WAE.

You can use static bypass 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). For information about how to centrally configure static bypass lists for a WAE or a group of WAEs, see the "Configuring Static Bypass Lists for WAEs" section on page 4-28.

Note We recommend that you use ACLs (redirect lists) on the WCCP-enabled router where possible, because that is the most efficient method to control traffic interception. However, when using inline interception, you can use static bypass lists or interception ACLs on the WAEs. 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 4-10. For information about how to configure an interception ACL for a WAE, see the "Configuring an Interception Access Control List" section on page 4-29.

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 6-33.

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 7-7.

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 the following types of device groups:

- Standard Device Group—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. There are two types of device groups:
  - Configuration Groups
  - Wide Area File Services (WAFS) Core Clusters (only used for WAFS legacy mode)

  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, which is a configuration group. This practice allows you to configure settings for the group, then apply those settings consistently across all your WAEs.

Note The AllDevicesGroup is a default device group that automatically contains all WAEs. In the AllDevicesGroup or any other device group, you should configure only the settings that you want to be consistent across all the WAEs in the group. Settings that apply to a single WAE should be configured on that device only and not on the device group.

- Baseline Group—A special type of device group used to configure a WAAS service consistently across multiple WAEs. There are three types of baseline groups:
  - File
  - Acceleration
  - Platform

  For example, if you want all your WAEs to contain the same set of application policies, we recommend that you create an Acceleration baseline group that contains all your custom and modified policies. When you assign WAEs to this group, the WAEs automatically inherit the application policies from the group. Anytime you need to change a policy, you make the change on the Acceleration baseline group and the change is propagated to the member devices. Setting up a baseline group is a way to apply consistent service settings across WAEs that reside in different device groups because WAEs can belong to separate device groups.

Note We recommend that you do not configure file and acceleration settings for a device group. Instead, use the File and Acceleration baseline groups for this purpose.

By default, WAAS Central Manager allows you to assign a device to multiple device groups (including baseline 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 WAFS policies.

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

- Prepositioning only works in a CIFS environment.

- 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.

PART 2

Installing and Configuring WAAS

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 appliances, WAE Network Modules (the NME-WAE family of devices), and SM-SRE modules running WAAS.

This chapter contains the following sections:

- About Device and Baseline Groups, page 3-1
- Working with Device Groups, page 3-2
- Working with Baseline Groups, page 3-10
- Working with Device Locations, page 3-14

About Device and Baseline 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.

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

You can organize your WAAS devices into the following types of device groups:

- Standard Device Group—A collection of devices that share common qualities and capabilities. Setting up groups based on their authentication settings as previously described is an example of a device group. There are two types of device groups: Configuration Group and WAFS Core Cluster. These are explained in more detail in the "Creating a New Device Group" section on page 3-3.
- Baseline Group—A special type of device group used to configure a WAAS service consistently across multiple devices. There are three types of baseline groups: File, Acceleration, and Platform.

  By default, all devices registered with the WAAS Central Manager are assigned to all three baseline groups.

  Baseline groups allow you to apply consistent service settings across devices that reside in different device groups.

  For example, if you have WAAS devices that reside in different device groups and you want all the devices to share the same application policies, you should make all your policy changes to the Acceleration baseline group. Whenever you create a new policy or modify an existing policy, those changes are distributed to each device that belongs to the Acceleration baseline group. If you make the policy changes to a specific device group, the devices that belong to your other groups are not updated with the policy changes.

Working with Device Groups

This section contains the following topics:

- Creating a Device Group, page 3-2
- Deleting a Device Group, page 3-6
- Viewing Device Group Assignments, page 3-6
- Viewing the Device Groups List, page 3-7
- Enabling or Disabling Device Group Overlap, page 3-7
- Overriding Group Configuration Settings, page 3-8
- Understanding the Impact of Assigning a Device to Multiple Device Groups, page 3-10

Creating a Device Group

This section contains the following topics:

- Creating a New Device Group, page 3-3
- Configuring the Settings for a Device Group, page 3-4
- Assigning Devices to a Configuration Device Group, page 3-5

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 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.
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 GUI navigation pane, choose **My WAN > Manage 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 From the Type drop-down list, choose one of the following options:

- **Configuration Group**—The standard type of device group.
- **Legacy WAFS Core Cluster**—A special type of device group that needs to be created only when configuring wide area file services (WAFS) using the legacy mode. For more information, see the "Configuring the Core Cluster" section on page 11-11.

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.

Do not check any of the Baseline check boxes. If you want to create a baseline group, see the "Customizing the Baseline Group Settings" section on page 3-11.

Note The system is configured with file, acceleration, and platform baseline groups by default. Only one of each of these types of baseline groups can be configured, so you won't be able to check any of the Baseline check boxes unless you have deleted one of these groups. Checking the Is Baseline check box will have no effect unless you are recreating one of these groups.

Step 6 (Optional) Enter comments about the group in the Comments field. The comments that you enter will appear in the Device Group window.

Step 7 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 8 (Optional) Customize the navigation pane 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 9 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, configure only the settings that you want to be consistent across all the devices. Settings that apply to a single device should be configured on that device only and not on the device group.

Note We recommend that you do not configure file and acceleration settings for a device group. Instead, use the File and Acceleration baseline groups for this purpose. For more information, see the "Working with Baseline Groups" section on page 3-10.

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

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Device Groups**.

The Device Groups window appears.

Step 2 Click the **Edit** icon next to the device group that you want to configure.

The Modifying Device Group window appears.

Step 3 Click the **Pages configured for this device group** arrow button to view which configuration windows have already been configured for the baseline 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 4 Customize the navigation pane 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 5 Use the navigation pane to navigate to each configuration window that you want to modify for this device group.

If a window 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 6 Make the necessary changes on the configuration 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 7 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 **My WAN > Manage Devices**, click the **Edit** icon next to the device that you want to assign to a group, and choose **Assign Groups** from the navigation pane. 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.

Note By default, all devices automatically join the AllDevicesGroup 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 AllDevicesGroup before you assign the device to another device group.

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

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Device Groups**.

Step 2 Click the **Edit** icon next to the device group to which you want to assign devices.

The Modifying Device Group window appears.

Step 3 In the navigation pane, choose **Assign Devices**.

The WAE Assignments window appears, displaying the WAEs devices assigned to various locations.

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 4 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 5 Click **Submit**.

A green check mark appears next to the assigned devices.

Step 6 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 WAEs** 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 GUI navigation pane, choose **My WAN > Manage Device Groups**. The Device Groups window appears.

Step 2 Click the **Edit** icon next to the name of the device group that you want to delete. The Modifying Device Group window appears.

Step 3 In the taskbar, click the **Delete Device Group** icon. You are prompted to confirm your decision to delete the device group.

Step 4 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 GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the device for which you want to view its group assignments.

The Device Dashboard window appears.

Step 3 In the Assignments section 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. 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 directly by browsing to the Assign Groups option in the navigation pane.

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

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Device Groups**.

Step 2 Click the **Edit** icon next to the group for which you want to view its device members.

The Modifying Device Group window appears.

Step 3 From the navigation pane, choose **Assign Devices**.

The WAE Assignments window appears, which shows all the 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 **My WAN > Manage Device Groups** in the WAAS Central Manager GUI.

This window displays the following information about each device group:

- Type of device group (either Configuration Group or WAFS Core Cluster).
- 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 (including baseline 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 GUI navigation pane, choose **Configure > 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 3-8
- Selecting Device Group Precedence, page 3-8
- Overriding the Device Group Settings on a Device, page 3-9

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 GUI navigation pane, choose **My WAN > Manage Device Groups**.

The Device Groups listing window appears.

Step 2 Click the **Edit** icon next to the device group with the settings that you want to force on all devices in that group.

The Modifying Device Group window appears.

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

The WAAS Central Manager GUI returns the following message:

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

Step 4 To force group settings across all devices in the device group, click **OK**.

Step 5 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 3-10.

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 GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the device that you want to set the device group precedence.

The Device Dashboard window appears.

Step 3 From the navigation pane, browse to the configuration window that contains the conflicting settings.

A drop-down list appears in the taskbar at the top of the window, as shown in Figure 3-1. 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.

Figure 3-1 Specifying the Device Group Precedence for a Configuration Window

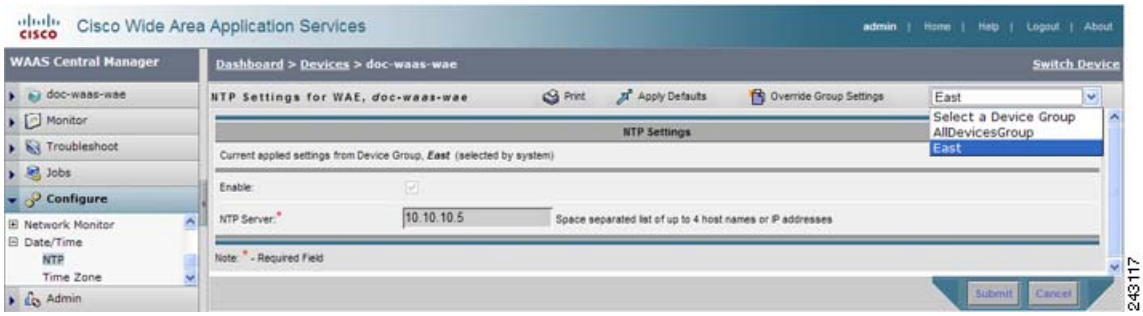

Step 4 From the drop-down list, choose the device group that you want this configuration page to inherit settings from, and click **Submit**.

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 GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the device that you want to override its group settings.

The Device Dashboard window appears.

Step 3 From the navigation pane, browse to the configuration window that contains the device group settings you want to override.

Step 4 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 5 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 6 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 3-8.

Working with Baseline Groups

A baseline group is a special type of device group used to configure a WAAS service consistently across multiple devices. The WAAS Central Manager GUI provides the following three types of baseline groups:

- File—Configures file services consistently across multiple devices.
- Acceleration—Configures the application policies consistently across multiple devices.
- Platform—Configures platform settings consistently across multiple devices.

For example, if you want all your devices to have the same set of application policies, we recommend that you create an Acceleration baseline group that contains all your custom and modified policies. When you assign all your devices to this group, the devices automatically inherit the application policies from the group. Anytime you need to change a policy, you make the change to the baseline group and the change is propagated to all your devices.

A device can be a member of multiple baseline groups. However, a particular service can have only one baseline group associated with it at any given time.

A baseline group is configured in the same way and functions in the same manner as a device group. You first create the baseline group, then configure or modify service settings for that group, then finally assign devices to the group.

This section contains the following topics:

- Configuring the Default Baseline Groups, page 3-11

- Switching the Baseline Group for a Service, page 3-13

Configuring the Default Baseline Groups

Table 3-2 describes the process for configuring the default baseline groups that come with your WAAS system.

Table 3-2 Checklist for Configuring the Default Baseline Groups

Task	Additional Information and Instructions
1. Customize the baseline group settings.	Changes the basic properties of the baseline group like whether newly activated devices automatically join the group. For more information, see the "Customizing the Baseline Group Settings" section on page 3-11.
2. Configure the service settings for the baseline group.	Configures the service settings that are unique to the baseline group. All devices that are a member of this group will automatically inherit these settings. For more information, see the "Configuring the Service Settings for a Baseline Group" section on page 3-12.
3. Assign devices to the baseline 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.

This section contains the following topics:

- Customizing the Baseline Group Settings, page 3-11
- Configuring the Service Settings for a Baseline Group, page 3-12

Customizing the Baseline Group Settings

The first step in configuring a baseline group is to customize the basic settings that determine the following:

- Whether all newly activated devices automatically join the baseline group.
- The configuration windows that you want to hide for the baseline group.

To customize the settings for a baseline group, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose one of the three baseline groups that you want to customize, as follows:

- For the File Services baseline group, choose **Configure > File Services > Baseline Group**.
- For the Acceleration Services baseline group, choose **Configure > Acceleration > Baseline Group**.
- For the Platform Services baseline group, choose **Configure > Platform > Baseline Group**.

The Selecting the (File, Acceleration, or Platform) Baseline Group window for the chosen service appears.

Note If this baseline group does not have a device group assigned to it, the Selecting the File Baseline Group window appears. From this window, select the device group that you want to assign to this baseline group, and then proceed with the rest of the instructions in this section.

Step 2 Change the name of the baseline group by entering the new name in the provided field.

The name must be unique and should be a name that is useful in distinguishing the baseline group from others on your system. The name cannot contain spaces or special characters.

Step 3 Check the **Automatically assign all newly activated devices to this group** check box to set this baseline group as the default device group for all newly activated devices. Only check this box if you want all your devices to inherit the settings of this baseline group.

Step 4 Enter your comments in the Comments field to provide a description of this baseline group.

The comments you enter will appear in the Device Group window.

Step 5 Click the **Pages configured for this device group** arrow to view the list of the windows in the WAAS Central Manager GUI that have been configured for this baseline group.

Step 6 Customize the navigation pane for this baseline 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 baseline group. Use this feature to remove from view any configuration windows that you do not need for this particular baseline group.

Step 7 Click **Submit**.

The windows that you selected to hide for this baseline group disappear from the navigation pane.

Step 8 Configure the service settings for this baseline group as described in the section that follows.

Configuring the Service Settings for a Baseline Group

Each baseline group should be configured to reflect the unique service settings that are not shared with the other baseline groups. For example, if you use the File baseline group to configure File Services on your devices and the Acceleration baseline group to configure the application policies, then these two baseline groups should be configured with different service settings. In this case, the File baseline group can be configured with Edge Server service enabled, and the Acceleration baseline group can be configured with custom or modified application policies.

To configure service settings for a baseline group, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose one of the three baseline groups that you want to configure, as follows:

- For the File Services baseline group, choose **Configure > File Services > Baseline Group**.
- For the Acceleration Services baseline group, choose **Configure > Acceleration > Baseline Group**.
- For the Platform Services baseline group, choose **Configure > Platform > Baseline Group**.

The Modifying Device Group window for the chosen service appears.

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

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

Step 3 Use the navigation pane to navigate to the configuration window that you want to modify for this baseline group.

Depending on the baseline group you are configuring, you will likely want to modify the following configuration windows:

- Configure > File Services—Configures service settings for the File baseline group.
- Configure > Acceleration—Configures service settings for the Acceleration baseline group.
- Configure > General Settings—Configures service settings for the Platform baseline group.

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

Step 4 After making the necessary changes to a configuration window, click **Submit**.

After you submit your changes, the modified configuration window is listed under Pages configured for this device group in the Modifying Device Group window.

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

Switching the Baseline Group for a Service

The WAAS Central Manager GUI allows you to switch the device group that is associated with a baseline group. When you switch a baseline group, you must choose a regular device group to take its place. During the switch, the regular device group that you choose is converted to a baseline group, and the baseline group that you remove is converted to a regular device group.

To remove a baseline group from a service and associate another baseline group in its place, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Configure > File** or **Acceleration** or **Platform > Baseline Group**.

The Modify Device Group window for the chosen service appears.

Step 2 Click the **Switch Group** icon in the taskbar. The WAAS Central Manager GUI returns the following message:

```
This action will remove this device group as the Baseline Group for this service. You can
then select another device group or create a new one to the Baseline Group for this
service. Do you wish to Continue?
```

Step 3 Click **OK** to remove the device group from the service.

The WAAS Central Manager GUI displays the Selecting the Baseline Group window for the chosen service.

Step 4 From the Select a Device Group to be the Baseline Group drop-down list, choose a device group, or choose the **Create New Device Group** option, as follows:

- If you choose a device group to be the baseline group for that service, the WAAS Central Manager GUI takes you to the Modify Device Group window for that device group.

- If you choose the Create New Device Group option, the WAAS Central Manager GUI takes you to the Create New Device Group window.

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 16-36.

You assign a device to a location when you activate the device as described in the "Modifying Device Properties" section on page 9-1.

You can work with locations by performing these tasks:

- Creating Locations, page 3-14
- Deleting Locations, page 3-15
- Viewing the Location Tree, page 3-15

Creating Locations

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

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage 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 9-1.

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 GUI navigation pane, choose **My WAN > Manage Locations**. The Locations window appears.

Step 2 Click the **Edit** icon next to the location that you want to delete.

The Modifying Location window appears.

Step 3 In the taskbar, click the **Delete Location** icon. You are asked to confirm your decision to delete the location.

Step 4 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 **My WAN > Manage Locations**. In the taskbar, click the **Location Trees** button. The location tree is displayed as shown in Figure 3-2.

Figure 3-2 Example of the Location Tree

CHAPTER 4

Configuring Traffic Interception

This chapter describes the WAAS software support for intercepting all TCP traffic in an IP-based network, based on the IP and TCP header information, and redirecting the traffic to wide area application engines (WAEs). This chapter describes the use of the Web Cache Communication Protocol (WCCP), policy-based routing (PBR), and inline mode for transparent redirection of traffic to WAEs.

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.

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 Configuration Fundamentals Configuration Guide* and the *Cisco IOS Configuration Fundamentals Command Reference*.

This chapter contains the following sections:

- Request Redirection Methods, page 4-1
- Using WCCP to Transparently Redirect TCP Traffic to WAEs, page 4-3
- Configuring Advanced WCCP Features on a WCCP-Enabled Router, page 4-7
- Managing WCCP Configurations for WAEs, page 4-12
- Configuring Egress Methods for Intercepted Connections, page 4-30
- Using Policy-Based Routing to Transparently Redirect All TCP Traffic to WAEs, page 4-33
- Using Inline Mode to Transparently Intercept TCP Traffic, page 4-42

Request Redirection 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 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, you can intercept traffic directly by using the inline mode with a WAE that has a Cisco WAE Inline Network Adapter.

In your WAAS network, traffic can be intercepted in these modes:

- Transparent mode (WCCP or PBR)
  - For application traffic, there are no configuration changes required to the client or the client-server applications. In promiscuous WCCP mode, application traffic is transparently redirected by network elements to the local WAE.

    Note The TCP promiscuous mode service (WCCP services 61 and 62) includes all protocols that use TCP as a transport, including the Common Internet File System (CIFS) protocol. When you enable the TCP promiscuous mode service on the routers and the WAEs, CIFS traffic is redirected to the WAEs because CIFS runs over TCP.

  - For CIFS traffic, the branch WAEs will accelerate the traffic based on the system configuration and policy. The WAEs do not advertise the names of file servers while operating in transparent mode unless a file server is configured for the disconnected mode of operation and the network becomes disconnected. CIFS traffic between clients and file servers relies on the client's ability to reach the server natively (through directed IP traffic or name resolution). With TCP promiscuous mode, a router redirects CIFS traffic (on TCP ports 139 or 445) to a local WAE, where it is optimized based on the local policy on that WAE. The only name service provided by the branch WAE in this mode is for local print services if local print services is configured.

    Note A branch WAE operates in only one of two modes: transparent (discussed above) or nontransparent (discussed below). This mode is configured on the WAE and applies to all file servers being accelerated by it. The configured mode is saved on the Central Manager and on the WAE.

- Nontransparent (explicit) mode (WCCP Version 2 disabled; applicable only to CIFS traffic when using the legacy file services mode)
  - For CIFS traffic, the Edge WAE publishes a file server name on the branch office network. This published name cannot be identical to the origin file server name due to NetBIOS name conflicts. Client computers must map drives from accelerated file servers using the published name as presented by the Edge WAE. This is the default mode.
  - For application traffic (non-CIFS), some form of interception is required in order to optimize the traffic. WCCP, PBR, inline mode, or CSM/ACE redirection must be configured; otherwise, non-CIFS traffic is unable to be optimized by WAAS.
- Inline mode

  The WAE physically and transparently intercepts traffic between the clients and the router. To use this mode, you must use a WAE with the Cisco WAE Inline Network Adapter installed.

Table 4-1 summarizes the transparent traffic interception methods that are supported in your WAAS network.

Table 4-1 Supported Methods of Transparent Traffic Interception

Method	Comment
WCCP Version 2	Used for transparent interception of application traffic and Wide Area File Services (WAFS) traffic. Used in branch offices and data centers to transparently redirect traffic to the local WAE. The traffic is transparently intercepted and redirected to the local WAE 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 "Using WCCP to Transparently Redirect TCP Traffic to WAEs" section on page 4-3.
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 L3 switch may be configured to use PBR to transparently intercept and route client and server traffic to WAE(s) 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 to Transparently Redirect All TCP Traffic to WAEs" section on page 4-33.
Inline	Used for transparent interception of application traffic and WAFS traffic. See the "Using Inline Mode to Transparently Intercept TCP Traffic" section on page 4-42.
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 WAE(s) 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 5-16.

Using WCCP to Transparently Redirect TCP Traffic to WAEs

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
- Multiple routers
- Multicasting of protocol messages between the WAE and the WCCP-enabled router
- Authentication of protocol packets
- Redirection of non-HTTP traffic
- Packet return (including GRE, allowing a WAE to reject a redirected packet and to return it to the router to be forwarded)
- Layer 2-caching (through router versus GRE) and 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). 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, interception ACLs, and static bypass.

The Cisco 2600, Cisco 2800, Cisco 3600, Cisco 3700, Cisco 3800, and Cisco 7600 series routers are supported, and can be manually configured and enabled with WCCP Version 2 support for use with the Cisco WAEs. The Catalyst 6000 and Catalyst 6500 series switches also support WCCP Version 2.

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, page 4-5
- Guidelines for File Server Access Methods, page 4-6

Guidelines for Configuring WCCP

When you configure transparent redirection on a WAE using WCCP Version 2, follow these general 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 4-11.
- 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 4-10.
- You must configure the WAE to accept redirected packets from one or more WCCP-enabled routers.
- You can quickly view a list of WCCP settings and services that you can configure on a WAE, from the WAAS CLI or the WAAS Central Manager GUI or the WAAS CLI. From the WAAS CLI, enter the **wccp** EXEC command followed by a question mark (?). The following sample output is from a WAE with WCCP Version 2 enabled:

```
WAE(config)# wccp ?
   access-list      Configure an IP access-list for inbound WCCP encapsulated traffic
   flow-redirect    Redirect moved flows
   router-list      Router List for use in WCCP services
   shutdown         Wccp Shutdown parameters
   tcp-promiscuous  TCP promiscuous mode service
   version          WCCP Version Number
```

- To configure basic WCCP, you must enable the WCCP service on at least one router in your network and on the WAE 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 (WCCP services 61 and 62) on the router and the WAE, as described in the *Cisco Wide Area Application Services Quick Configuration Guide.*

- In order for the WAE to function in TCP promiscuous mode, the WAE uses WCCP Version 2 services 61 and 62. 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 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. In the configuration example provided in the *Cisco Wide Area Application Services Quick Configuration Guide*, the CLI is used 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 "Managing WCCP Configurations for WAEs" section on page 4-12.
- 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 a WCCP-Enabled Router" section on page 4-7.
- To ensure consistency, we recommend that you configure WCCP settings on a device group basis instead of on an individual device, and the device group should contain WAEs only from a single WCCP service farm. It is best not to configure WCCP using a device group that contains WAEs that belong to multiple farms, since WCCP settings may need to be different in different farms.
- 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 defaults to the assignment method that the router supports, even if the WAE is configured differently, unless you configure the WAE with the strict assignment method option (by checking the Only Use Selected Assignment Method check box). When using the strict assignment method option, the WAE only joins the farm if the assignment method configured on the WAE is supported by the router.
- A WAE will join 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 will reconfigure and the WAE will be removed from the farm.

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 a WCCP-Enabled Router

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:

- Configuring a Router to Support WCCP Service Groups, page 4-7
- Configuring IP Access Lists on a Router, page 4-10
- Setting a Service Group Password on a Router, page 4-11
- Configuring a Loopback Interface on the Router, page 4-12
- Configuring Router QoS for WCCP Control Packets, page 4-12

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.*

Configuring a Router to Support WCCP Service Groups

WCCP Version 2 enables a set of branch WAEs in a WAE 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 4-1.

Figure 4-1 Service Groups with WCCP Version 2

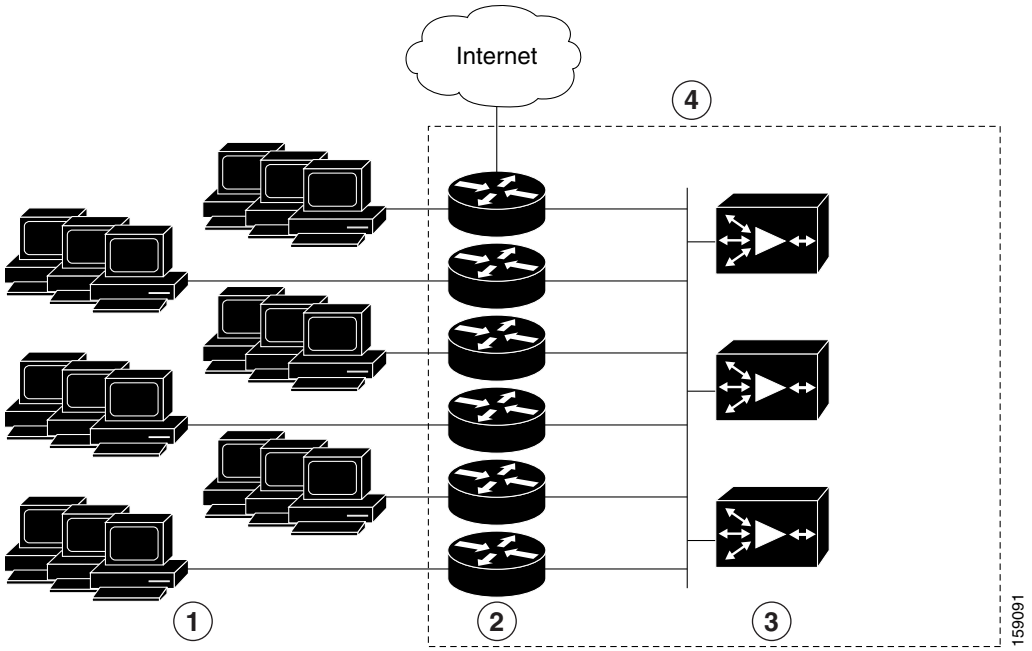

1	Clients requesting file services	**3**	Branch WAEs
2	Cisco routers	**4**	WAE service group

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 GUI (the "Viewing or Modifying WCCP Service Masks for WAEs" section on page 4-23) 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).

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.

ip wccp {**web-cache** | *service-number*} [**group-address** *groupaddress*]

The following example shows how to enable the TCP promiscuous mode service (WCCP Version 2 services 61 and 62) on a router in a group of routers that have a group address of 224.10.10.1 in a multicast deployment:

```
Router(config)# ip wccp 61 group-address 224.10.10.1
Router(config)# ip wccp 62 group-address 224.10.10.1
```

For a unicast deployment, enable TCP promiscuous mode service as follows:

```
Router(config)# ip wccp 61
Router(config)# ip wccp 62
```

On each WAE in a multicast deployment, configure only the multicast address in the WCCP router list, as follows:

```
WAE(config)# wccp router-list 1 224.10.10.1
```

On each WAE in a unicast deployment, configure multiple unicast router addresses in the WCCP router list, one for each router in the service group.

Additionally, in a multicast deployment, you need to configure each router to accept multicast packets on one interface, using commands similar to the following:

```
Router(config)# interface vlan817
Router(config-subif)# ip wccp 61 group-listen
Router(config-subif)# ip wccp 62 group-listen
Router(config-subif)# ip pim dense-mode
```

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. Should the administrator disable flow protection, 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 an Interception Access Control List" section on page 4-29. For information about how to configure a static bypass list, see the "Configuring Static Bypass Lists for WAEs" section on page 4-28. You can also configure interface ACLs on WAEs to control management access to the WAE, as described in Chapter 8, "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 to be redirected to a different WAE in the farm, negatively impacting performance. Affected traffic includes management traffic, legacy CIFS tunnel traffic, and so on. 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 software 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 (WCCP Version 2 services 61 and 62) on the router:

```
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 GUI to specify the service group password on a WAE (or device group), see the "Viewing or Modifying the WCCP Settings on WAEs" section on page 4-18.

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

Beginning with version 4.2.1, WAAS sends WCCP control packets marked with a differentiated services code point (DSCP) value of 192. (Previously, 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.

Managing WCCP Configurations for WAEs

This section contains the following topics:

- Load Balancing and WAEs, page 4-13
- Packet-Forwarding Methods, page 4-15
- WCCP Flow Redirection on WAEs, page 4-18
- Viewing or Modifying the WCCP Settings on WAEs, page 4-18
- Creating a WCCP Service Mask for an Existing WCCP Service, page 4-23
- Viewing or Modifying WCCP Service Masks for WAEs, page 4-23
- Viewing a WCCP Router List Configuration for WAEs, page 4-24
- Modifying the Configuration of WCCP Router Lists for WAEs, page 4-25
- Deleting a WCCP Router List from WAEs, page 4-25
- Defining Additional WCCP Router Lists on WAEs, page 4-26
- Configuring WAEs for a Graceful Shutdown of WCCP, page 4-27
- Configuring Static Bypass Lists for WAEs, page 4-28

- Configuring an Interception Access Control List, page 4-29

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*.

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.

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.

The term *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.

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 4-2.)

Figure 4-2 Load Balancing Through Hashing of IP Addresses

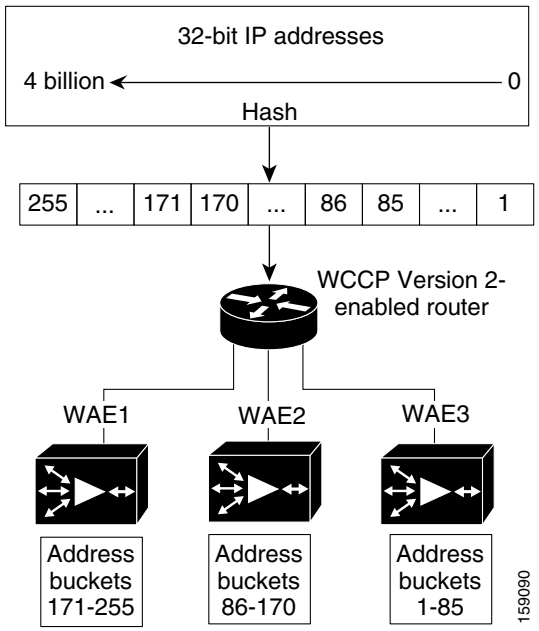

Note Packets that the WAEs do not service are tunneled back to the same router from which they were received. When a router receives a formerly redirected packet, it knows not to redirect it again.

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 load unevenly, however, 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 (flow protection makes sure that it is the same WAE throughout the session).

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 IOS version 12.4(20)T or later.

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 router(s) 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 datacenter WAE, to improve 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.

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 4-2 describes the packet-forwarding methods.

Table 4-2 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, page 4-16
- Layer 3 GRE as a Packet-Forwarding Method, page 4-17
- Layer 2 Redirection as a Packet-Forwarding Method, page 4-17

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.
- CIFS packets are received by the branch WAE destined to a server that is not configured to be cached by the branch WAE (only in WAFS legacy mode).
- 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 "Modifying the Configuration of WCCP Router Lists for WAEs" section on page 4-25.

Cisco Express Forwarding (CEF) is required for WCCP and must be enabled on the router.

WCCP also allows you to configure multiple routers (router lists) 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 generic routing encapsulation (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.

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 enabled on a WAE. Flow protection reduces the impact on existing client TCP connections when branch WAEs are added and removed from a service group. The client impact is reduced because of flow protection in the following situations:

- 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.

Viewing or Modifying the WCCP Settings on WAEs

To ensure consistency, we recommend that you configure the WCCP settings on a device group basis instead of on an individual device, and the device group should contain WAEs only from a single WCCP service farm.

Note If you add a WAE with a Cisco WAE Inline Network Adapter to a device group that has WCCP configured, the WCCP device group settings are not automatically applied to the inline WAE. If you want to configure the inline WAE to use WCCP, you must manually force the WCCP device group settings to be applied.

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 (WCCP Version 2 services 61 and 62) as described in the *Cisco Wide Area Application Services Quick Configuration Guide.*

To modify the WCCP settings for a WAE (or group of WAEs), follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or device group) for which you want to modify the WCCP settings or services.

Step 3 In the navigation pane, choose **Configure > Interception > WCCP > Settings**. The WCCP Settings window appears. (See Figure 4-3.)

Figure 4-3 Modifying the WCCP Settings

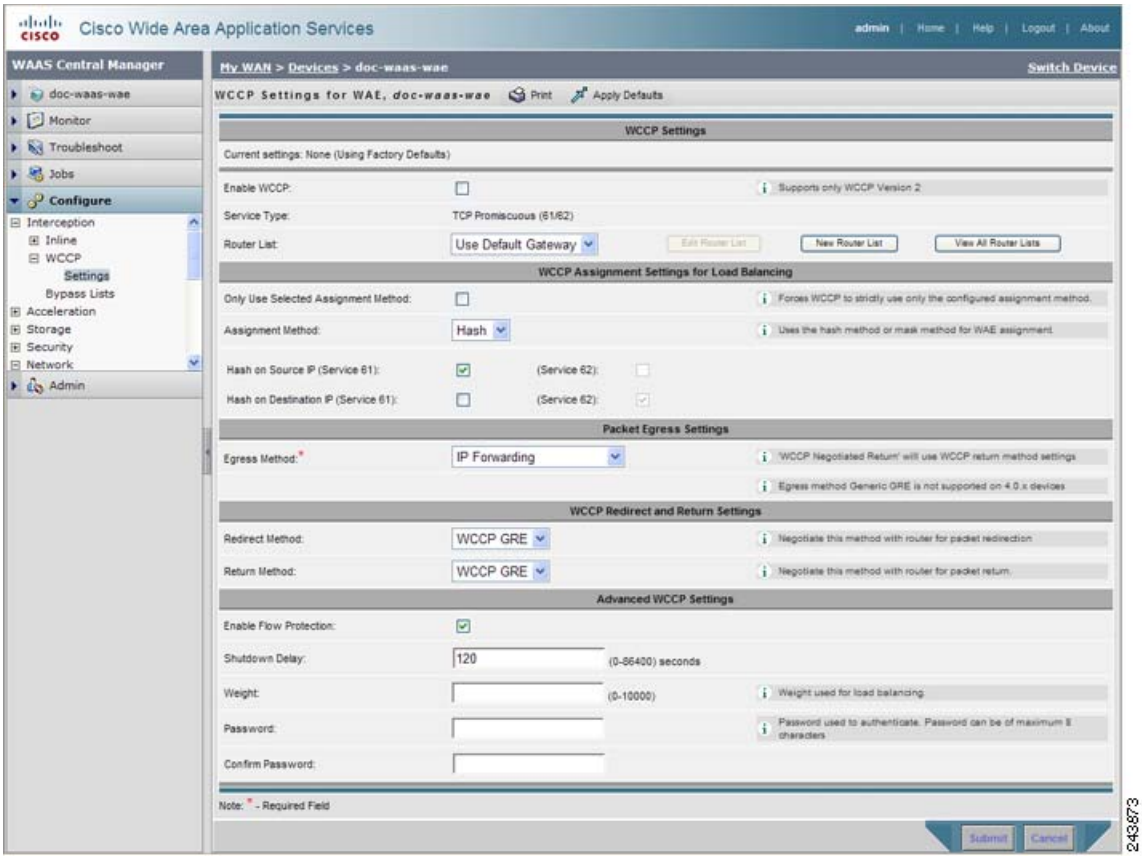

Step 4 Check the current settings for the chosen device (or device group):

- To keep the current settings and to close the window, click **Cancel**.
- To modify the current settings, change the current setting as described in the rest of this procedure.

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 5 Check the **Enable WCCP** check box to enable WCCP Version 2 on the chosen device (or device group), or uncheck the check box to disable WCCP on the chosen device (or device groups).

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.

Step 6 Associate a router list with the WCCP TCP promiscuous mode service by choosing the appropriate number of the WCCP router list from the Router List drop-down list. Alternatively, you can choose the default, which is Use Default Gateway. The default router list contains the single IP address of the default gateway of the WAE device.

Only configured WCCP router lists are displayed in the drop-down list. As part of the initial configuration of your WAAS network, you will have already created at least one WCCP router list for your branch WAE and a second WCCP router list for your data center WAE, as described in the *Cisco Wide Area Application Services Quick Configuration Guide*. For more information about WCCP router lists, see the following sections:

- Modifying the Configuration of WCCP Router Lists for WAEs, page 4-25
- Deleting a WCCP Router List from WAEs, page 4-25
- Defining Additional WCCP Router Lists on WAEs, page 4-26

Step 7 (Optional) Modify the current load-balancing settings in the WCCP Assignment Settings for Load Balancing area as follows:

a. To force WCCP to use the configured assignment method only, check the **Only Use Selected Assignment Method** check box. When applied, choose either of the two load-balancing methods: hash assignment or mask assignment.

You can specify only one load-balancing method (hashing or masking) per WCCP service in a branch WAE group.

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.

b. From the Assignment Method drop-down list, choose the type of WAE load-balancing assignment method to use (for more information, see the "Load Balancing and WAEs" section on page 4-13):

- Choose **Hash** (the default) to use the hash method. Then follow Steps 8 and 9 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. Then skip to Step 10 to define or edit a service mask.

Step 8 (Optional) To define the load-balancing hash for WCCP service 61 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 9 (Optional) To define the load-balancing hash for WCCP service 61on 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.

Note For more information about load balancing, see the "Load Balancing and WAEs" section on page 4-13.

Step 10 (Optional) Click the **Edit Mask** button to modify an existing service mask for the WCCP service. For more information about modifying service masks, see the "Viewing or Modifying WCCP Service Masks for WAEs" section on page 4-23. (This button is shown only if the mask assignment method is used.)

Step 11 (Optional) Click the **View Masks Configured for All Services** button to view a list of all configured WCCP service masks for the service. (This button is shown only if the mask assignment method is used.) The WCCP Service Mask Settings windows appears.

Step 12 From the WCCP Service Mask Settings window, you can perform the following tasks:

- To edit a WCCP service mask, click the **Edit WCCP Service Mask** icon next to the service mask that you want to modify. The Modifying WCCP Service Mask window appears. (See Figure 4-4.)

Figure 4-4 Modifying a WCCP Service Mask

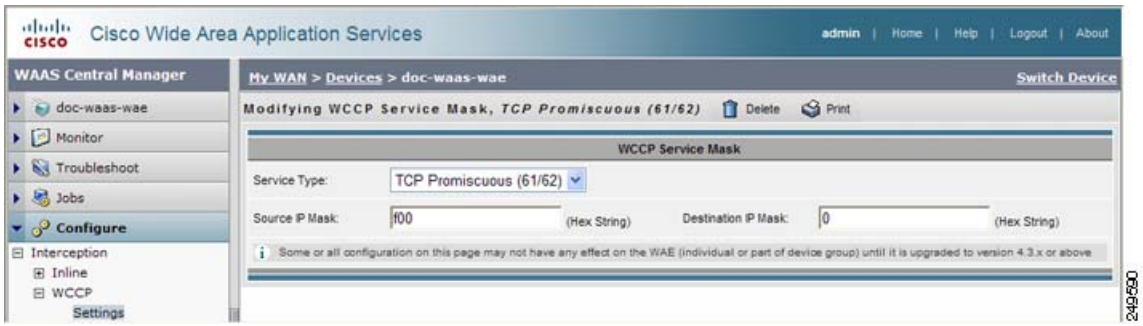

Change the values of the settings that you want to modify and click **Submit**, as follows:

- In the Source IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, 0xFE000000) used to match the packet source IP address. The range is 0x00000000–0xFE000000. The default is 0xF00.

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.

- In the Destination IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, 0xFE000000) used to match the packet destination IP address. The range is 0x0000000–0XFE000000. The default is 0x00000000.

Note You can also edit a service mask by clicking the **Edit Mask** button in the Modifying WCCP Service window. (See Figure 4-3.)

- To delete an existing WCCP service mask, click the **Edit WCCP Service Mask** icon next to the service mask that you want to delete. The Modifying WCCP Service Mask window appears. (See Figure 4-4.) Click the **Delete WCCP Service Mask** icon in the taskbar and click **Submit**.

Step 13 (Optional) In the Packet Egress Settings area, from the Egress Method drop-down list, choose the type of method to use to return optimized packets to the router or switch: **IP Forwarding** (the default), **WCCP Negotiated Return**, or **Generic GRE**. For more details on choosing the egress method, see the "Configuring Egress Methods for Intercepted Connections" section on page 4-30.

Step 14 (Optional) Modify the current packet redirect and return method settings in the WCCP Redirect and Return Settings area as follows:

a. From the Redirect Method drop-down list, choose the type of packet redirection (forwarding) method to use:

- Choose **WCCP GRE** (the default) to use Layer 3 GRE packet redirection.
- Choose **WCCP L2** to permit the WAE (or device group) 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.

WCCP on a router or switch can take advantage of switching hardware that either partially or fully implements the traffic interception and redirection functions of WCCP in hardware at Layer 2. The WAE can then perform a Layer 2 or MAC address rewrite redirection if it is directly connected to a compatible Cisco switch. This redirection processing is accelerated in the switching hardware, which makes this method a more efficient method than Layer 3 redirection using GRE. The WAE must have a Layer 2 connection with the router or switch. Because there is no requirement for a GRE tunnel between the switch and the WAE, the switch can use a cut-through method of forwarding encapsulated packets if you choose L2. For more information, see the "Packet-Forwarding Methods" section on page 4-15.

b. From the Return Method drop-down list, choose the type of method to use to return nonoptimized (bypassed) packets to the router:

- Choose **WCCP GRE** (the default) to use GRE packet return.
- Choose **WCCP L2** to use Layer 2 rewriting for packet return.

Step 15 (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 (or device groups) when they come up or are reassigned new traffic. For more information, see the "WCCP Flow Redirection on WAEs" section on page 4-18. Flow protection is enabled by default.

b. In the Shutdown Delay field, specify the maximum amount of time (in seconds) that the chosen device (or device group) 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 Version 2.

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 8 characters in length. 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 4-11.

To configure WCCP settings from the CLI, you can use the **wccp flow-redirect**, **wccp router-list**, **wccp shutdown**, **wccp tcp-promiscuous**, and **wccp version** 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 4-27.

Creating a WCCP Service Mask for an Existing WCCP Service

To create a service mask for an existing WCCP service for a WAE (or group of WAEs), follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Devices Group**).

Step 2 Click the **Edit** icon next to the device (or device group) for which you want to create a WCCP service mask.

Step 3 In the navigation pane, choose **Configure > Interception > WCCP > Settings**. The WCCP Settings window appears.

Step 4 From the Assignment Method drop-down list, choose Mask.

Step 5 Click the **Create Mask** button. The Creating New WCCP Service Mask window appears.

Step 6 In the Source IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, 0xFE000000) used to match the packet source IP address. The range is 0x00000000–0xFE000000. The default is 0xF00.

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.

Step 7 In the Destination IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, 0xFE000000) used to match the packet destination IP address. The range is 0x0000000–0XFE000000. The default is 0x00000000.

Step 8 Click **Submit** to save the settings for the WCCP service mask.

Viewing or Modifying WCCP Service Masks for WAEs

To view or modify a service mask for a WCCP service that is configured for a WAE (or group of WAEs), follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Devices Group**).

Step 2 Click the **Edit** icon next to the device (or device group) for which you want to view or modify a WCCP service mask.

Step 3 In the navigation pane, choose **Configure > Interception > WCCP > Settings**. The WCCP Settings window appears.

Step 4 From the Assignment Method drop-down list, choose Mask.

Step 5 Click the **Edit Mask** button. The Modifying WCCP Service Mask window appears.

Step 6 In the Source IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, 0xFE000000) used to match the packet source IP address. The range is 0x00000000–0xFE000000. The default is 0xF00.

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.

Step 7 In the Destination IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, 0xFE000000) used to match the packet destination IP address. The range is 0x0000000–0XFE000000. The default is 0x00000000.

Step 8 Click **Submit** to save the new settings for the WCCP service mask.

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 will 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.

Viewing a WCCP Router List Configuration for WAEs

To centrally view the list of currently defined WCCP router list for a WAE (or group of WAEs), follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or device group) for which you want to display a WCCP router list.

Step 3 In the navigation pane, choose **Configure > Interception > WCCP > Settings**. The WCCP Settings window appears.

Step 4 Click the **View All Router Lists** button.

The WCCP Router List Configurations window for the chosen device (or device group) appears.

The configuration for the WCCP router lists (the number of the router list and IP addresses of each router that is included in each router list) is displayed.

Note To modify the configuration of a specific WCCP router list, click the **Edit** icon next to the router list and use the displayed Modifying Router List to modify the chosen router list. For more information about modifying router lists, see the "Modifying the Configuration of WCCP Router Lists for WAEs" section on page 4-25. For information about how to delete a WCCP router list from a WAE (or group of WAEs), see the "Deleting a WCCP Router List from WAEs" section on page 4-25.

To view a router list from the CLI, you can use the **show wccp routers** EXEC command.

Modifying the Configuration of WCCP Router Lists for WAEs

To centrally modify the configuration of a WCCP router list (for example, add or delete a router from a router list) for a WAE (or group of WAEs), follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or device group) for which you want to modify the router list configuration.

Step 3 In the navigation pane, choose **Configure > Interception > WCCP > Settings**.

The WCCP Settings window appears.

Step 4 Click the **Edit Router List** button. The Modifying WCCP Router List window appears.

Step 5 Add a router to the chosen router list by entering the router IP address in the Add Router field and clicking the **Add Router** button.

Step 6 Remove a router from the chosen router list by checking the check box next to the IP address of the router that you want to remove and clicking the **Remove Router** button.

Step 7 Click **Submit** to save the settings.

Deleting a WCCP Router List from WAEs

When you delete a router list, the WCCP Version 2 services that have been configured to use this router list are also deleted. Ensure that the WCCP service is associated with a different router list, if required, before deleting the previously configured router list.

To centrally delete a WCCP router list (for example, add or delete an IP address from a router list) for a WAE (or group of WAEs), follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or device group) for which you want to delete a WCCP router list.

Step 3 In the navigation pane, choose **Configure > Interception > WCCP > Settings**. The WCCP Settings window appears.

Step 4 Click the **Edit Router List** button. The Modifying WCCP Router List window appears.

Step 5 Remove all of the listed routers from the chosen router list by checking the check the check box next to the IP address of the router that you want to remove and clicking the **Remove Router** button.

Step 6 After you have removed all the routers from the chosen router list (for example, router list 2), click the **Delete Router List** icon in the taskbar.

The system displays a dialog box asking you to confirm that you want to permanently delete the router list configuration. To confirm your decision, click **OK**. The selected router list and the associated WCCP services are deleted from the chosen device (or device group).

Defining Additional WCCP Router Lists on WAEs

As part of configuring a WCCP service on a WAE, you must create a list of WCCP Version 2-enabled routers that support the TCP promiscuous service for the WAE. You can define a WCCP router list through the WAAS CLI (the **wccp router-list** global configuration command) or the WAAS Central Manager GUI.

Typically, WAAS administrators will use the WAAS CLI to define their initial set of WCCP router lists, as described in the *Cisco Wide Area Application Services Quick Configuration Guide.* After you have used the WAAS CLI to complete the initial configuration of your WCCP router lists, we recommend that you use the WAAS Central Manager GUI to centrally manage and modify your WCCP router list configurations for your WAEs.

Each WAE includes a default router list (router list 8) that is configured with the IP address of the router that is defined as the default gateway for the WAE.

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 (WCCP Version 2 services 61 and 62) as described in the *Cisco Wide Area Application Services Quick Configuration Guide*.

To centrally define additional WCCP router list for a WAE (or group of WAEs), follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or device group) for which you want to create a WCCP router list.

Step 3 In the navigation pane, choose **Configure > Interception > WCCP > Settings**. The WCCP Settings window appears.

Step 4 Click the **New Router List** button.

The Creating New WCCP Router List window appears. (See Figure 4-5.)

Figure 4-5 Creating a New WCCP Router List Example Screen

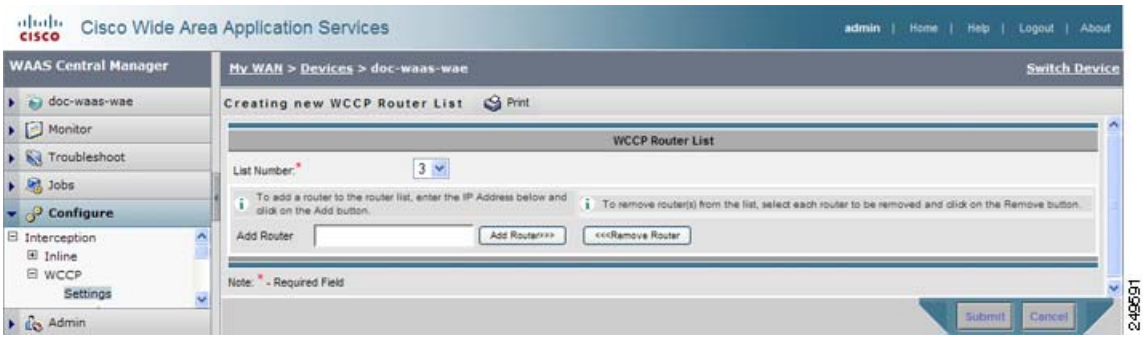

In this example, the number 3 is preselected in the List Number drop-down list because there are already two WCCP router lists defined for the chosen device (or device group). Router list 1 has already been defined for the WCCP router in the data center that will be transparently redirecting traffic to the data center WAE, and router list 2 was defined for the WCCP router in the branch office that will be transparently redirecting traffic to the branch WAE that resides in the same branch office.

Step 5 In the Add Router field, specify the IP address of the router to be added to router list 3.

You must enter at least one IP address. All IP addresses added must be unique within the router list. Otherwise, an error message is displayed on submit.

Step 6 Click **Add** to add an IP address to router list 3.

This list represents the IP address of every WCCP router that is to transparently redirect traffic to the chosen WAE (or group of WAEs) for the TCP promiscuous mode service.

The window refreshes and the addresses are listed in numerical order. The order might not match the order in which IP addresses were entered.

Step 7 Click **Submit** to save the router list or to save any edits you have made to the router IP addresses.

To define a router list from the CLI, you can use the **wccp router-list** global configuration command.

After you create a WCCP router list on a WAE or group of WAEs, ensure that WCCP Version 2 is enabled and configured on the WCCP routers that are included in this new router list as described in the *Cisco Wide Area Application Services Quick Configuration Guide*.

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.

The WAAS Central Manager GUI allows you to disable WCCP Version 2 on a WAE. You can also perform this task locally through the CLI (by entering the **no wccp version** CLI command on the WAE).

To disable WCCP for a chosen device or device group, uncheck the **Enable WCCP** check box in the WAAS Central Manager's WCCP Settings window. (See Figure 4-3.)

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 an 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

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 4-10. For information about how to configure an interception ACL for a WAE, see the "Configuring an Interception Access Control List" section on page 4-29.

To centrally configure a static bypass list for a WAE (or group of WAEs), follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or device group) for which you want to create a static bypass list.

Step 3 From the navigation pane, choose **Configure > Interception > Bypass Lists**.

Step 4 In the taskbar, click the **Create New WCCP/Inline Bypass List** icon. The Creating new WCCP/Inline Bypass List window appears.

Step 5 Enter the IP address for the client in the Client Address field.

Step 6 Enter the IP address for the server in the Server Address field.

Step 7 Check **Submit** to save the settings.

To configure a static bypass list from the CLI, you can use the **bypass static** global configuration command.

Aggregating Static Bypass Lists

An individual WAE device can have static bypass lists defined and can belong to device groups that have other static bypass lists defined.

In the WCCP Bypass Lists window, the Aggregate Settings radio button controls how static bypass lists are aggregated for an individual device, as follows:

- Choose **Yes** if you want to configure the device with all static bypass lists 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 static bypass lists 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 an Interception Access Control List

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 an interception ACL on the branch WAE, 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 branch WAE. 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.

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.

To use an interception ACL, first define an ACL (see Chapter 8, "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 a WAE, follow these steps:

Step 1 Follow the instructions in Chapter 8, "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 GUI navigation pane, choose **My WAN > Manage Devices**.

Step 3 Click the **Edit** icon next to the name of the device for which you want to use an interception ACL.

Step 4 From the navigation pane, choose **Configure > Interception > Interception Access List**.

Step 5 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 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 6 Check **Submit** to save the settings.

To configure an interception ACL from the CLI, you can use the **ip access-list** and **interception access-list** global configuration commands.

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.

Configuring Egress Methods for Intercepted Connections

The WAAS software supports the following egress methods for WCCP intercepted connections:

- IP forwarding
- WCCP GRE return
- Generic GRE

The default egress method is IP forwarding. If you do not configure the egress method, then the WAE uses 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. 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 4-32.

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.

The WCCP bypass traffic uses WCCP GRE as the return method, regardless of the CLI configuration.

The WCCP negotiated return and generic GRE egress methods do not apply to the inline mode of operation.

To configure the egress method for WCCP intercepted connections from the Central Manager GUI, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or device group) for which you want to configure the egress method.

Step 3 In the navigation pane, choose **Configure > Interception > WCCP > Settings**. The WCCP Settings window appears (see Figure 4-3).

Step 4 In the Packet Egress Settings area, from the Egress Method drop-down list, choose **IP Forwarding**, **WCCP Negotiated Return**, or **Generic GRE**.

Step 5 Click **Submit**.

To configure the interception and egress method for WCCP GRE packet return from the CLI, use the **egress-method** global configuration command:

```
WAE(config)# egress-method negotiated-return intercept-method wccp
```

To configure the interception and egress method for IP forwarding from the CLI, use the **egress-method** global configuration command:

```
WAE(config)# egress-method ip-forwarding intercept-method wccp
```

To configure the interception and egress method for the generic GRE egress method from the CLI, configure an intercepting router list and configure the generic GRE egress method, as follows:

```
WAE(config)# wccp router-list 1 192.168.68.98
WAE(config)# egress-method generic-gre intercept-method wccp
```

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 4-32.

To view the egress method that is configured and that is being used on a particular WAE, use the **show egress-methods** 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. Use the command **ip wccp redirect exclude in**.

This section contains the following topics:

- Multipoint Tunnel Configuration, page 4-32
- Point-To-Point Tunnel Configuration, page 4-33

Multipoint Tunnel Configuration

Consider a deployment in which there are two intercepting routers and two WAEs in the farm. Each WAE's configuration would look like the following:

```
wccp router-list 1 192.168.1.1 192.168.2.1
wccp tcp-promiscuous router-list-num-1
egress-method generic-gre intercept-method wccp
```

Each router can configure a single multipoint GRE tunnel to the WAE farm.

The router 1 configuration would look like the following:

```
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 Provisioning an IP address for the tunnel interface is what enables it for IP, allowing 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 to Transparently Redirect All TCP Traffic to WAEs

Policy-based routing (PBR), introduced in the Cisco IOS Software 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 4-39.

Note For a description of the PBR commands that are referenced in this section, see the *Cisco Quality of Service Solutions Command Reference*.

Figure 4-6 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 4-6 Example of Using PBR or WCCP Version 2 for Transparent Redirection of All TCP Traffic to WAEs

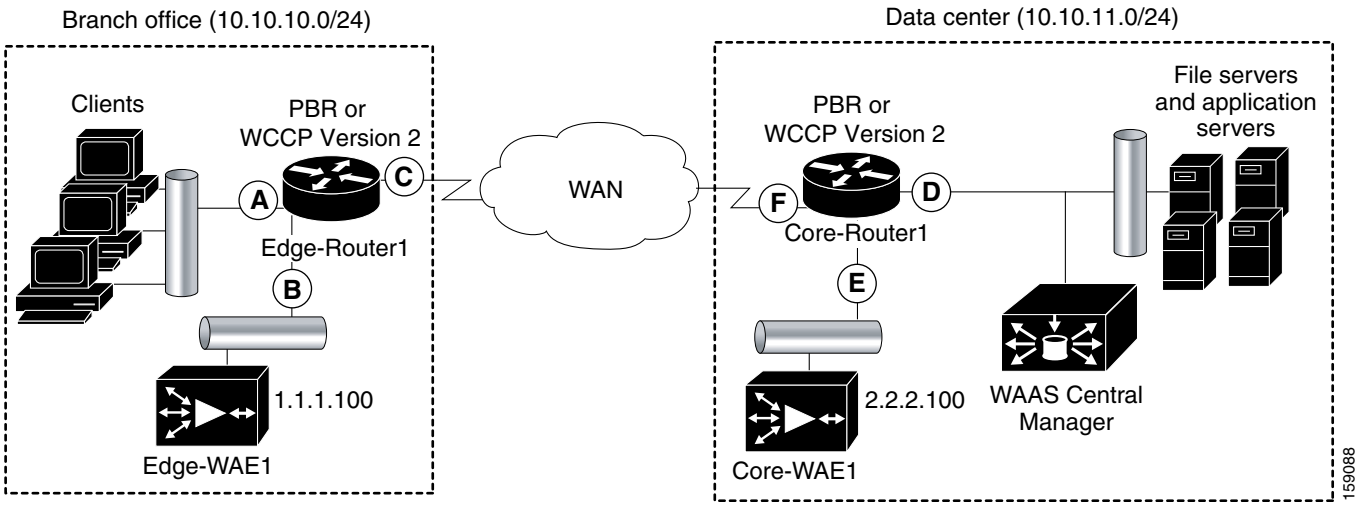

Table 4-3 provides a summary of the router interfaces that you must configure to use PBR or WCCP Version 2 to transparently redirect traffic to a WAE.

Table 4-3 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 4-6, redundancy (for example, redundant routers, switches, WAEs, WAAS Central Managers, and routers) is not depicted.

The following example 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 4-6).

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.

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 4-40.
- 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 4-40.
- 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 4-41.

Note In this section, the term 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 to Transparently Intercept TCP Traffic

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 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 4-7. 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* located on Cisco.com.

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 using hardware accelerated WCCP interception or load balancing with the Cisco Application Control Engine (ACE).

Figure 4-7 Inline Interception

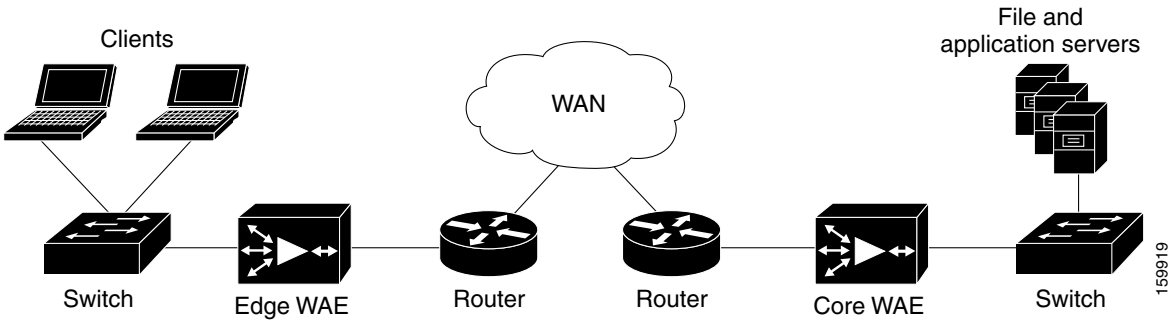

Note Inline mode and WCCP redirection are exclusive. You cannot configure inline mode if the WAE is configured for WCCP operation. Inline mode is the default mode when the Cisco WAE Inline Network Adapter is installed in a WAE device. If you add a WAE with a Cisco WAE Inline Network Adapter to a device group that has WCCP configured, the WCCP device group settings are not automatically applied to the inline WAE. If you want to configure the inline WAE to use WCCP, you must manually force the WCCP device group settings to be applied.

Note A WAE with a Cisco WAE Inline Network Adapter can be configured as a Central Manager, but the inline interception functionality will not be available.

The Cisco WAE Inline Network Adapter contains two or four Ethernet ports. On adapters that have four Ethernet ports, the ports are grouped into two logical groups. Each group has one LAN-facing port and one WAN-facing port. Typically, you use just one group, and connect the LAN-facing port to a switch and the WAN-facing port to a router. On adapters with four ports, the second group of interfaces is provided if you are using a network topology where you need to connect the WAE to two routers. Traffic that enters on one interface in a group exits the device on the other interface in the same group.

The WAE-7341/7371/674 support two installed Cisco WAE Inline Network Adapters, providing a total of eight inline Ethernet ports.

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 4-47.

Traffic that flows through the Cisco WAE Inline Network Adapter 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, where all traffic is mechanically bridged between the LAN and WAN interfaces in each group. The Cisco WAE Inline Network Adapter also operates in bypass mode when the WAE is powered off or starting up. Additionally, you can manually put the Cisco WAE Inline Network Adapter into bypass mode.

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 with the Cisco WAE Inline Network Adapter installed to provide higher availability if a device fails. For details, see the "Clustering Inline WAEs" section on page 4-50.

Note When a WAE that has a Cisco WAE Inline Network Adapter installed enters bypass mode, the switch and router ports to which it is connected may have to reinitialize, and this 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.

This section contains the following topics:

- Configuring Inline Interface Settings, page 4-44
- Configuring an IP Address on an Inline Interface, page 4-47
- Configuring VLANs for Inline Support, page 4-49
- Clustering Inline WAEs, page 4-50

Configuring Inline Interface Settings

To configure inline interface settings, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**. (You cannot configure inline interface settings from Device Groups.)

The Devices window appears, listing all the devices configured in the WAAS network.

Step 2 Click the **Edit** icon next to the device for which you want to modify the inline settings.

The Device Dashboard window appears.

Step 3 In the navigation pane, choose **Configure > Interception > Inline > Inline Interfaces**.

The Inline Interfaces window appears, listing the inline interface groups available on the device. Click the **Edit Inline Interface** icon next to the inline interface group that you want to modify.

The Modifying Inline Interface window appears, displaying the inline interface configurations for a particular slot and group. (See Figure 4-8.)

Figure 4-8 Modifying Inline Interface Window

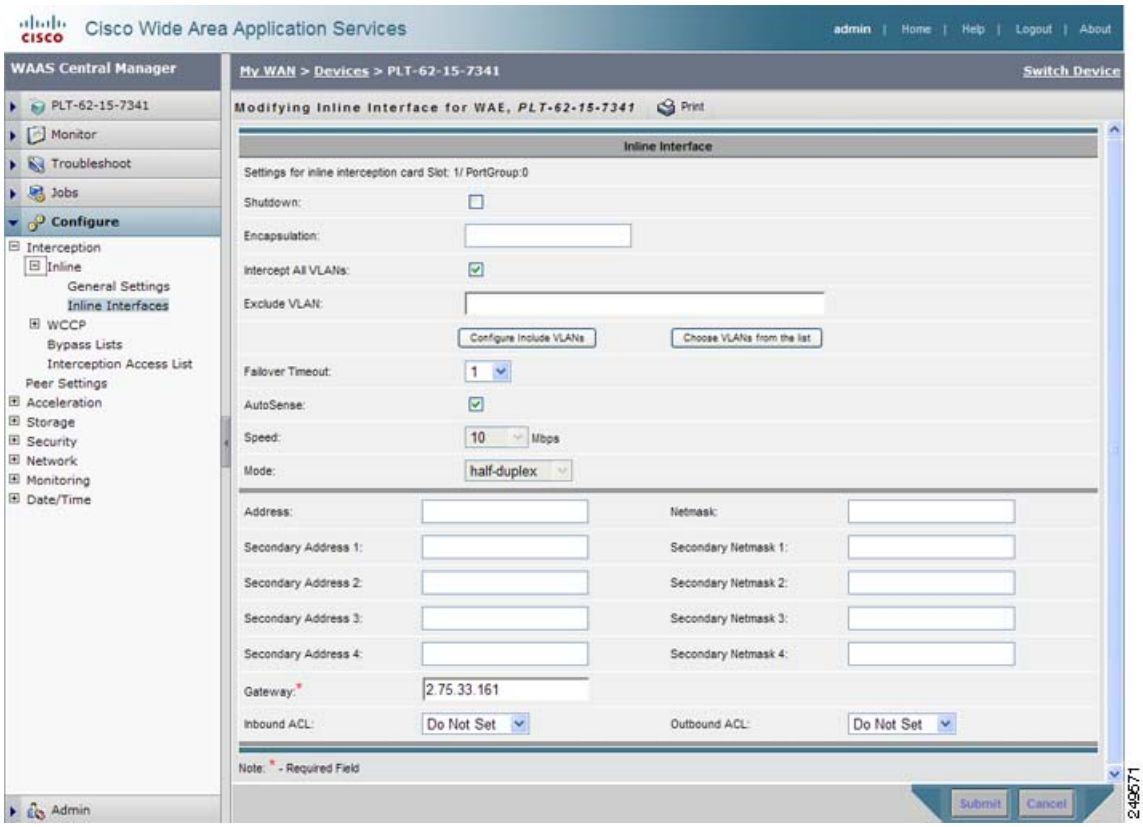

Step 4 Check the **Shutdown** check box to shut down the interface group. This setting bridges traffic across the LAN/WAN interfaces without any processing.

Step 5 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 4-47.

Step 6 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.

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 the Cisco WAE Inline Network Adapter is installed in a WAE device.

Step 7 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 pick 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 4-49.

c. Choose the VLAN ranges to include or exclude by doing either of the following:

- Click next to each VLAN range that you want to include for optimization on this inline interface group. The icon changes to . All VLANs that are not included for optimization are excluded.
- Click next to each VLAN range that you want to exclude from optimization on this inline interface group. The icon changes to .
- Click in the taskbar to select all available VLAN ranges for optimization, or click in the taskbar to exclude all VLAN ranges for optimization.

d. Click **Submit**.

Step 8 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.

Step 9 Configure the Speed and Mode port settings as follows:

a. Uncheck the **AutoSense** check box, which is enabled by default.

b. From the Speed drop-down list, choose a transmission speed (**10**, **100**, or **1000** Mbps).

c. From the Mode drop-down list, choose a transmission mode (**full-duplex** or **half-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 In the Address field, enter an IP address for the inline interface, if you want to assign an IP address.

Step 11 In the Netmask field, enter a subnet mask for the inline interface.

Step 12 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 13 In the Gateway field, enter the default gateway IP address.

Step 14 (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 15 (Optional) From the Outbound ACL drop-down list, choose an IP ACL to apply to outbound packets.

Step 16 Click **Submit**.

To configure inline interception from the CLI, use the **interface InlineGroup** global configuration command.

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.

To configure VLAN ID checking, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**. (You cannot configure network settings from Device Groups.)

The Devices window appears, listing all the devices configured in the WAAS network.

Step 2 Click the **Edit** icon next to the device for which you want to modify the network settings.

The Device Dashboard window appears.

Step 3 In the navigation pane, choose **Configure > Interception > Inline > General Settings**.

The General Settings window appears, listing the VLAN ID connection checking control.

Step 4 Click the **VLAN ID Connection Check** check box to enable VLAN ID connection checking. Uncheck the box to disable it. The default setting is enabled.

Step 5 Click **Submit**.

To configure VLAN ID checking from the CLI, use the **inline vlan-id-connection-check** global configuration command.

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" section on page 4-44.

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 traffic IP address in the *Device Name* > Activation window, in the Management IP field.

If a WAE with the Cisco WAE Inline Network Adapter 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 page at **Configure > Interception > Inline > Inline Interfaces** (see the "Configuring Inline Interface Settings" section on page 4-44).

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, port channel, and Cisco Discovery Protocol (CDP).
- 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.
- If you have configured a WAE to use the inline interfaces for all traffic and have not assigned IP addresses to the built-in interfaces, and then you downgrade the WAE to a version prior to 4.0.15, internal services on the WAE that are using the IP address assigned to the inline interface will lose connectivity. Before downgrading the WAE, configure the built-in interfaces with the appropriate network settings.
- Configuring IP addresses on the inline interfaces in a WAE is not supported if the WAAS Central Manager is running a software version prior to 4.0.15, even if the WAE is running version 4.0.15 or later.
- If you downgrade the WAAS Central Manager to a version prior to 4.0.15, in a deployment where a WAE is configured to use the inline interfaces for all traffic, the WAAS Central Manager will lose communication with the WAE because it cannot recognize the inline interface as the primary

interface. To avoid losing communication, either configure the WAE to accept management traffic on one of the built-in interfaces, or set the inline interface IP address as the alternate management address in the Device Name > Activation window, in the Management IP field.

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 GUI navigation pane, choose **Configure > Platform > Vlans**.

The Vlans window appears, listing 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. (See Figure 4-9.)

Figure 4-9 Creating New Vlan Window Example

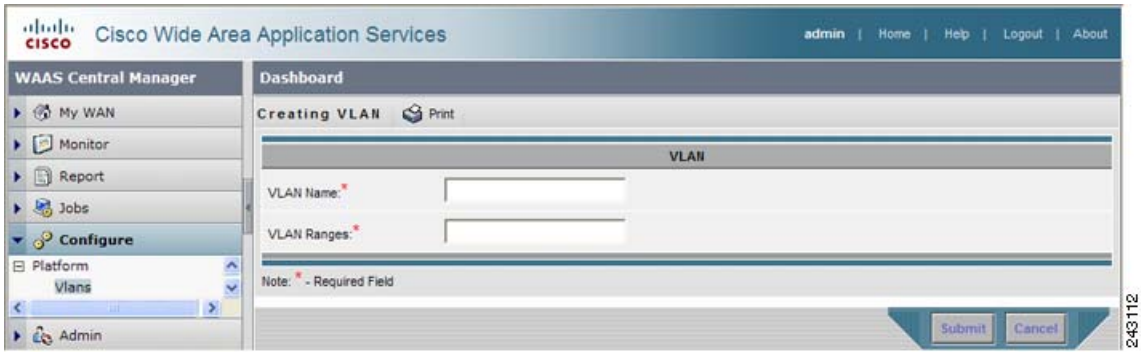

Step 3 In the VLAN Name field, enter a name for the VLAN list.

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" section on page 4-44. 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" section on page 4-44.

Clustering Inline WAEs

You can serially cluster two WAE devices with the Cisco WAE Inline Network Adapter installed 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 the 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 Cisco WAE Inline Network Adapter is connected to the LAN port of the next Cisco WAE Inline Network Adapter, as shown in Figure 4-10.

Figure 4-10 Inline Cluster

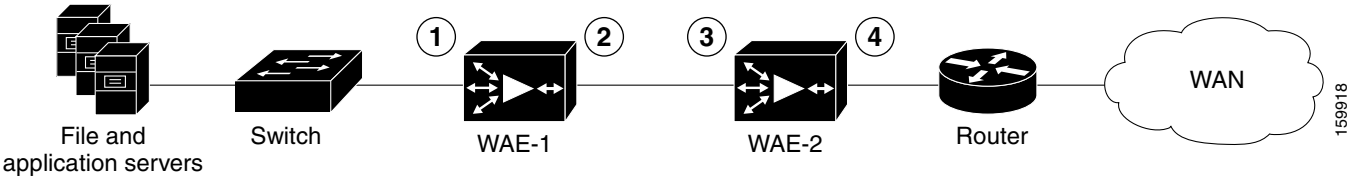

1	Inline LAN port on WAE-1	**3**	Inline LAN port on WAE-2
2	Inline WAN port on WAE-1	**4**	Inline WAN port on WAE-2

In a serial cluster, all traffic between the switch and router passes through all inline WAEs. In Figure 4-10, 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.

To disable peer optimization between WAEs in a serial cluster, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**. (You cannot configure peer settings from Device Groups.)

The Devices window appears, listing all the devices configured in the WAAS network.

Step 2 Click the **Edit** icon next to the device for which you want to configure the peer optimization settings. You can choose either WAE from the two that are paired as a serial cluster.

The Device Dashboard window appears.

Step 3 In the navigation pane, choose **Configure > Peer Settings**.

The Peer Settings window appears. (See Figure 4-11.)

Figure 4-11 Configuring Peer Settings

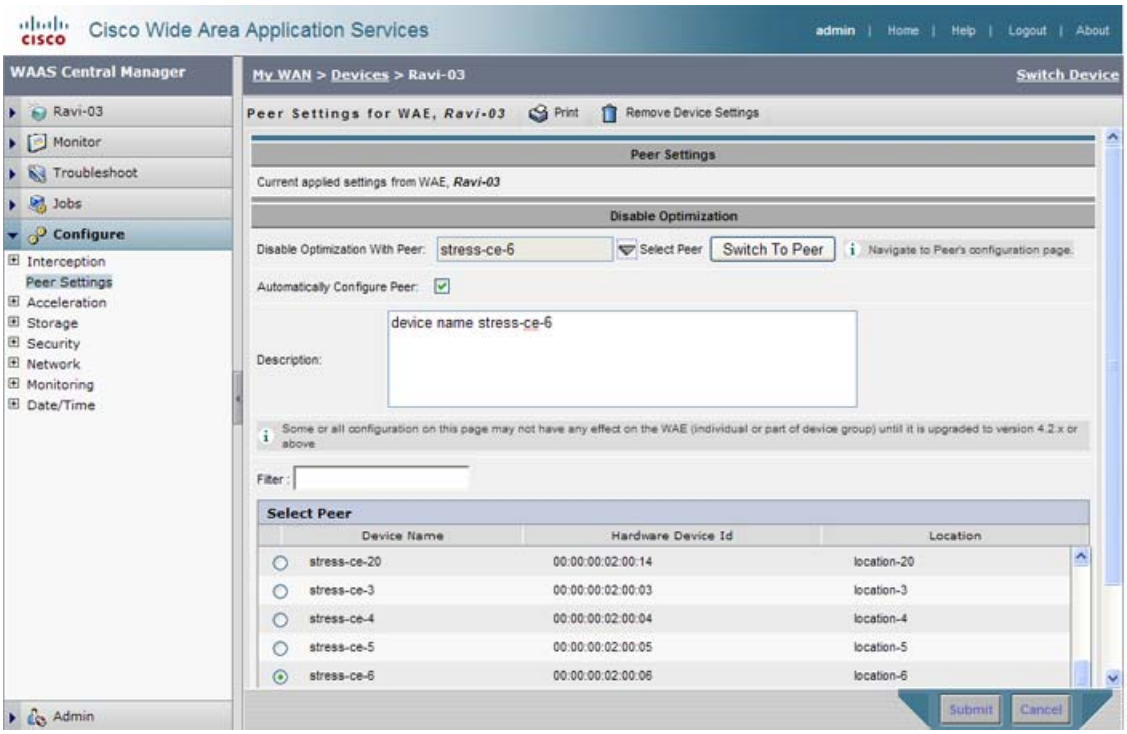

Step 4 Click the Select Peer triangle control to display in the lower part of the window other WAEs that are registered with this Central Manager (see the Select Peer area).

Step 5 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 6 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 7 In the Description field, enter a description for the peer. The default description is the device name of the peer.

Step 8 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 GUI navigation pane, choose **Configure > Peer Settings**. The Peer Settings window appears, as shown in Figure 4-12.

Figure 4-12 Viewing Peer Settings Status

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.

CHAPTER 5

Configuring Network Settings

This chapter describes how to configure basic network settings such as creating additional network interfaces to support network traffic, specifying a DNS server, enabling 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:

- Configuring Network Interfaces, page 5-1
- Configuring a Load-Balancing Method for Interfaces, page 5-9
- Configuring TCP Settings, page 5-9
- Configuring Static IP Routes, page 5-13
- Configuring CDP Settings, page 5-14
- Configuring the DNS Server, page 5-14
- Configuring Windows Name Services, page 5-15
- Configuring Directed Mode, page 5-16

Configuring Network Interfaces

During initial setup, you chose an initial interface and either configured it for DHCP or gave it a static IP address. This section describes how to configure additional interfaces using options for redundancy, load balancing, and performance optimization.

This section contains the following topics:

- Configuring a Standby Interface, page 5-2
- Configuring the Primary Standby Interface, page 5-3
- Configuring Multiple IP Addresses on a Single Interface, page 5-5
- Modifying Gigabit Ethernet Interface Settings, page 5-5
- Configuring Port-Channel Settings, page 5-7

- Configuring Interfaces for DHCP, page 5-8

We recommend that you use the WAAS Central Manager GUI 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**.

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.

Configuring a Standby Interface

In this procedure, you configure a logical interface called a standby interface. After you set up the parameters for this logical interface, you must associate physical interfaces with the standby interface to create the standby group. (A standby group consists of two or more physical interfaces.) In the WAAS Central Manager GUI, you create the standby group by joining physical interfaces to the standby group and assigning one physical interface as primary. (See the ""Configuring the Primary Standby Interface" section on page 5-3.")

Standby interfaces remain unused unless an active interface fails. When an active network interface fails (because of cable trouble, Layer 2 switch failure, or other failure), and that interface is part of a standby group, a standby interface can start to carry traffic and take the load off the failed interface. With standby interface configuration, only one interface is in use at a given time.

To configure standby interfaces, you must assign each physical interface to a standby group. The following operating considerations apply to standby groups:

- A standby group consists of two or more physical interfaces.
- The maximum number of standby groups on a WAAS device is one.
- 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.
- You cannot assign an IP address to both a standby group and a port channel. Only one virtual 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.

To configure a standby interface, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**. The Devices window appears.

Step 2 Click the **Edit** icon next to the device for which you want to configure a standby interface. The Device Dashboard window appears.

Step 3 In the navigation pane, choose **Configure > Network > Network Interfaces**. The Network Interfaces window for the device appears. (See Figure 5-1.)

Figure 5-1 Network Interfaces for Device Window

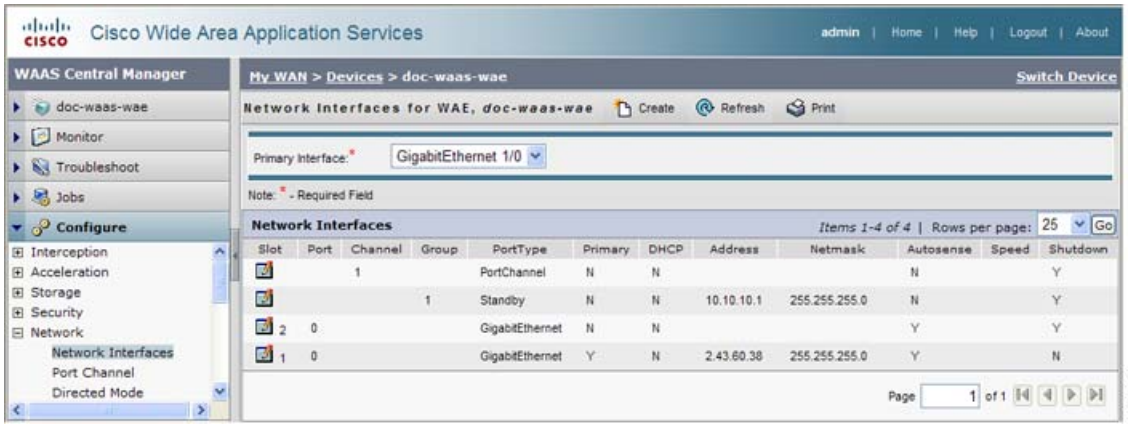

Step 4 In the taskbar, click the **Create New Interface** icon. The Creating New Network Interface window appears.

Step 5 From the Port Type drop-down list, choose **Standby**. The window refreshes with fields for configuring the standby group settings.

Step 6 In the Address field, specify the IP address of the standby group.

Step 7 In the Netmask field, specify the netmask of the standby group.

Step 8 Check the **Shutdown** check box to shut down the hardware interface. By default, this option is disabled.

Step 9 In the Gateway field enter the default gateway IP address. If an interface is configured for DHCP, then this field is read only.

Step 10 Click **Submit**.

Step 11 Configure the interface priority setting as described in the ""Configuring the Primary Standby Interface" section on page 5-3."

Configuring the Primary Standby Interface

After you have configured a logical standby interface using the WAAS Central Manager GUI, you configure the standby group by joining 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 5-2.")

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 GUI, choose **My WAN > Manage Devices**. The Devices window appears.

Step 2 Click the **Edit** icon next to the device for which you want to configure a standby interface. The Device Dashboard window appears.

Step 3 In the navigation pane, choose **Configure > Network > Network Interfaces**. The Network Interfaces window for the device appears.

Step 4 Click the **Edit** icon next to the physical interface that you want to join to a standby group. The Modifying Network Interface window appears. (See Figure 5-2.)

Do not choose a logical interface (standby or portchannel) in this step.

Figure 5-2 Modifying Network Interface Window

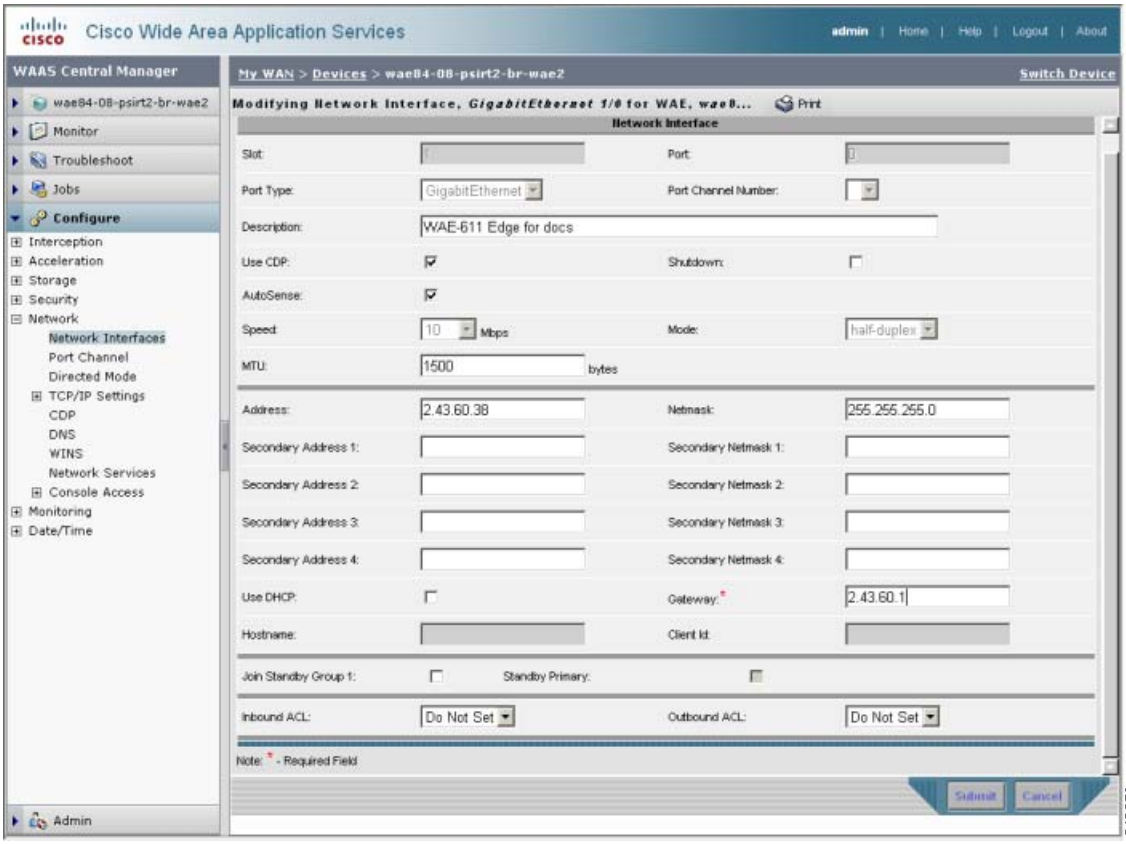

Step 5 Complete the following steps to join the interface to a standby group and specify it as the primary standby interface:

a. Check the **Join Standby Group 1** check box.

b. (Optional) Check the Standby Primary check box if you want the interface to be the primary (active) interface in the standby group.

Step 6 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 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.

To configure multiple IP addresses on a single interface, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**. The Devices window appears.

Step 2 Click the **Edit** icon next to the device for which you want to configure interface settings. The Device Dashboard window appears.

Step 3 In the navigation pane, choose **Configure > Network > Network Interfaces**. The Network Interfaces listing window appears.

Step 4 Click the **Edit** icon for the GigabitEthernet physical interface that you want to modify. The Modifying Network Interface window appears.

Note Do not choose a logical interface (standby or port channel) in this step. You cannot configure multiple interfaces on a logical interface.

Step 5 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 6 Click **Submit**.

Modifying Gigabit Ethernet Interface Settings

To modify the settings of an existing Gigabit Ethernet interface, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

The Devices window appears, listing all the device types configured in the WAAS network.

Step 2 Click the **Edit** icon next to the device for which you want to modify the interface settings.

The Device Dashboard window appears.

Step 3 In the navigation pane, choose **Configure > Network > Network Interfaces**.

The Network Interfaces window appears, listing the network interfaces configured on particular slots and ports.

Note On an NME-WAE device, 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*.

Step 4 Click the **Edit Network Interface** icon next to the Gigabit Ethernet interface that you want to modify.

The Modifying Network Interface window appears, displaying the interface configurations on a particular slot and port. (See Figure 5-2.)

Note Some of the fields in the window are not available. Interface configurations for slot, port, and port type are set for physical interfaces during initial startup or by using the WAAS CLI.

Note When configuring the internal interface (GigabitEthernet 1/0) on an NME-WAE 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 details, see the document *Configuring Cisco WAAS Network Modules for Cisco Access Routers*.

Step 5 Check the **Use CDP** check box to enable 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 5-14.

Step 6 Check the **Shutdown** check box to shut down the hardware interface.

Step 7 Check the **AutoSense** check box to set the interface to autonegotiate the speed and mode.

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 8 Manually configure the interface transmission speed and mode settings as follows:

a. Uncheck the **AutoSense** check box.

b. From the Speed drop-down list, choose a transmission speed (**10**, **100**, or **1000** Mbps).

c. From the Mode drop-down list, choose a transmission mode (**full-duplex** or **half-duplex**).

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 9 Specify a value (in bytes) in the MTU field to set the interface Maximum Transmission Unit (MTU) size.

The range is 88–1500 bytes. The MTU is the largest size of IP datagram that can be transferred using a specific data link connection.

Step 10 Enter a new IP address in the Address field to change the interface IP address.

Step 11 Enter a new netmask in the Netmask field to change the interface netmask.

Step 12 Click **Submit**.

Configuring Port-Channel Settings

WAAS software supports the grouping of up to four same-speed network interfaces into one virtual interface. This grouping capability allows you to set or remove a virtual interface that consists of two Gigabit Ethernet interfaces. 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.

To configure port-channel settings, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**. The Devices window appears.

Step 2 Click the **Edit** icon next to the name of the device for which you want to configure interfaces. The Device Dashboard window appears.

Step 3 In the navigation Pane, choose **Configure > Network > Network Interfaces**. The Network Interfaces window appears, listing all the interfaces for the chosen device.

Step 4 In the taskbar, click the **Create New Interface** icon. The Creating New Network Interface window appears.

Step 5 From the Port Type drop-down list, choose **Port Channel**.

The window refreshes and provides fields for configuring the network interface settings.

Step 6 In the Port Channel Number drop-down list, choose **1** for the port-channel interface number. (Only one port channel is supported.)

Step 7 Check the **Shutdown** check box to shut down this interface. By default, this option is disabled.

Step 8 In the Gateway field, enter the default gateway IP address.

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 (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 12 (Optional) From the Outbound ACL drop-down list, choose an IP ACL to apply to outbound packets.

Step 13 Click **Submit**.

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 virtual interface can be configured with an IP address.

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

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 GUI navigation pane, choose **My WAN > Manage Devices**. The Devices window appears.

Step 2 Click the **Edit** icon next to the name of the device for which you want to configure interface settings. The Device Dashboard window appears.

Step 3 In the navigation pane, choose **Configure > Network > Network Interfaces**. The Network Interfaces listing window appears.

Step 4 Click the **Edit** icon for the GigabitEthernet physical interface that you want to modify. The Modifying Network Interface window appears.

Note Do not choose a logical interface (standby or port channel) 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 device, because this interface can be configured only through the host router CLI. For details, see the document *Configuring Cisco WAAS Network Modules for Cisco Access Routers*.

Step 5 Scroll down the window and check the **Use DHCP** check box.

When this check box is checked, the secondary IP address and netmask fields are disabled.

Step 6 In the Hostname field, specify the hostname for the WAAS device or other device.

Step 7 In the Client Id field, specify the configured client identifier for the device.

The DHCP server uses that identifier when the WAAS device requests the network information for the device.

Step 8 Click **Submit**.

Configuring a Load-Balancing Method for 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 5-7.

To configure load balancing, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group with the port channel that you want to configure for load balancing.

Step 3 In the navigation Pane, choose **Configure > Network > Port Channel**.

Step 4 From the Load Balancing Method drop-down list, choose a load-balancing method:

- **round robin**—Round robin allows traffic to be distributed evenly among all interfaces in the channel group. The other balancing options give you the flexibility to choose specific interfaces (by IP address) when sending an Ethernet frame. This option is selected by default.
- **src-dst-ip-port**—The distribution function is based on a combination of source and destination IP addresses/ports. This load-balancing method has replaced the dst-ip method on devices running WAAS version 4.1.3 or later.

Step 5 Click **Submit**.

To configure a load-balancing method from the CLI, you can use the **port-channel** 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 in tuning these parameters. In nearly all environments, the default TCP settings are adequate. Fine tuning of 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the WAAS device or device group for which you want to configure TCP settings. The Device Dashboard window appears.

Step 3 In the navigation pane, choose **Configure > Network > TCP/IP Settings > TCP/IP**. The TCP/IP Settings window appears.

Step 4 Make the necessary changes to the TCP settings.

See Table 5-1 for a description of each TCP field in this window.

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 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 5-1 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 disabled. For more information, see the "Explicit Congestion Notification" section on page 5-11.
Initial Send Congestion Window Size	Initial congestion window size value in segments. The range is 1 to 10 segments. The default is 2 segments. For more information, see the "Congestion Windows" section on page 5-11.
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 5-11. **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 5-12.

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, page 5-11
- Congestion Windows, page 5-11
- Retransmit Time Multiplier, page 5-11
- TCP Slow Start, page 5-12
- Path MTU Discovery, page 5-12

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 5-12.")

You can modify the sender's retransmit timer by using the Retransmit Time Multiplier field in the WAAS Central Manager GUI. 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 5-11.") 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Devices Group**).

Step 2 Click the **Edit** icon next to the device or device group that you want to configure.

Step 3 In the navigation pane, choose **Configure > Network > TCP/IP Settings > Static Routes**. The IP Route Entries window appears.

Step 4 In the taskbar, click the **Create New IP Route Entry** icon. The Creating New IP Route window appears.

Step 5 In the Destination Network Address field, enter the destination network IP address.

Step 6 In the Netmask field, enter the destination host netmask.

Step 7 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 8 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

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group that you want to configure.

Step 3 From the navigation pane, choose **Configure > Network > CDP**. The CDP Settings window appears.

Step 4 Check the **Enable** check box to enable CDP support. By default, this option is enabled.

Step 5 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 6 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 7 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group that you want to configure.

Step 3 From the navigation pane, choose **Configure > Network > DNS**. The DNS Settings window appears.

Step 4 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 5 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 6 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group for which you want to configure Windows name services.

Step 3 From the navigation pane, choose **Configure > Network > WINS**. The Windows Name Services Settings window appears.

Step 4 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 5 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 6 In the WINS server field, enter the hostname or IP address of the Windows Internet Naming Service (WINS) server.

Step 7 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 IOS 12.4(11)T2 and later) or the **inspect waas** command (for FWSM 3.2(1) and later and PIX 7.2(3) and later).

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 Cisco WAE Inline Network Adapters) with routable, non-NATed IP addresses. When using directed mode with inline mode, you must configure the Cisco WAE Inline Network Adapter 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 GUI or 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or device group) for which you want to configure directed mode.

Step 3 In the navigation pane, choose **Configure > Network > Directed Mode**. The Directed Mode Settings window appears. (See Figure 5-3.)

Figure 5-3 Directed Mode Settings Window

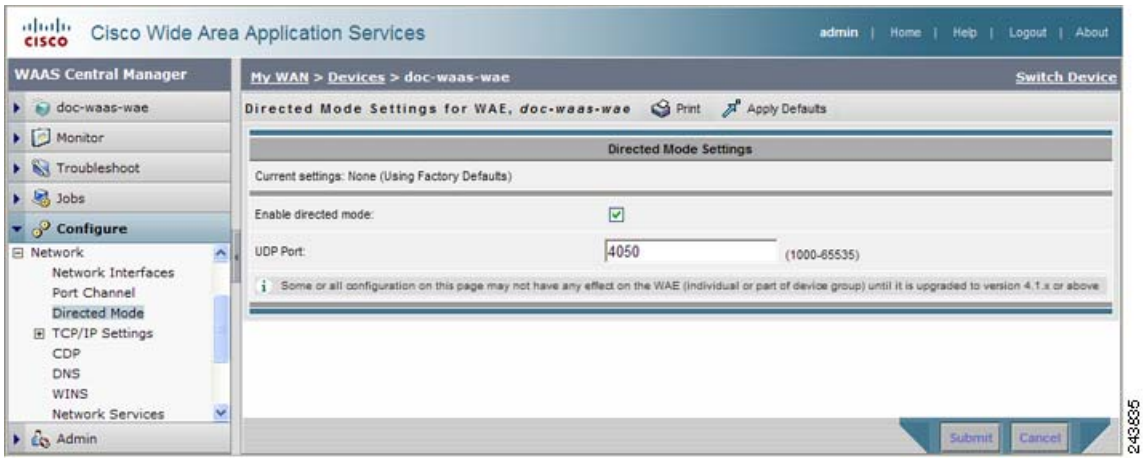

Step 4 Check the **Enable directed mode** check box to activate directed mode.

Step 5 Enter a port number in the UDP Port field to configure a custom UDP port for directed mode. The default is port 4050.

Step 6 Click **Submit** to save the settings.

To configure directed mode from the CLI, use the **directed-mode** global configuration command.

CHAPTER 6

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:

- About Administrative Login Authentication and Authorization, page 6-1
- Configuring Administrative Login Authentication and Authorization, page 6-6
- Configuring AAA Command Authorization, page 6-32
- Configuring AAA Accounting for WAAS Devices, page 6-33
- Viewing Audit Trail Logs, page 6-34

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 7, "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 6-33.

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.

If the predefined password for this superuser account has not been changed on a WAAS Central Manager device, the following dialog box is displayed each time you use the account to log in to the WAAS Central Manager GUI. (See Figure 6-1.)

Figure 6-1 Message Indicating the Predefined Password for the Superuser Account Should Be Changed

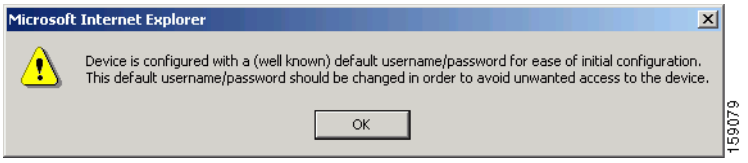

If you have not changed the predefined password for this superuser account, the console will also display the following message each time you use the account to log in to the WAAS CLI on a WAAS device:

```
Device is configured with a (well known) default username/password
for ease of initial configuration. This default username/password
should be changed in order to avoid unwanted access to the device.
```

```
System Initialization Finished.
waas-cm#
```

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 7-7.

Figure 6-2 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 6-2 Authentication Databases and a WAE

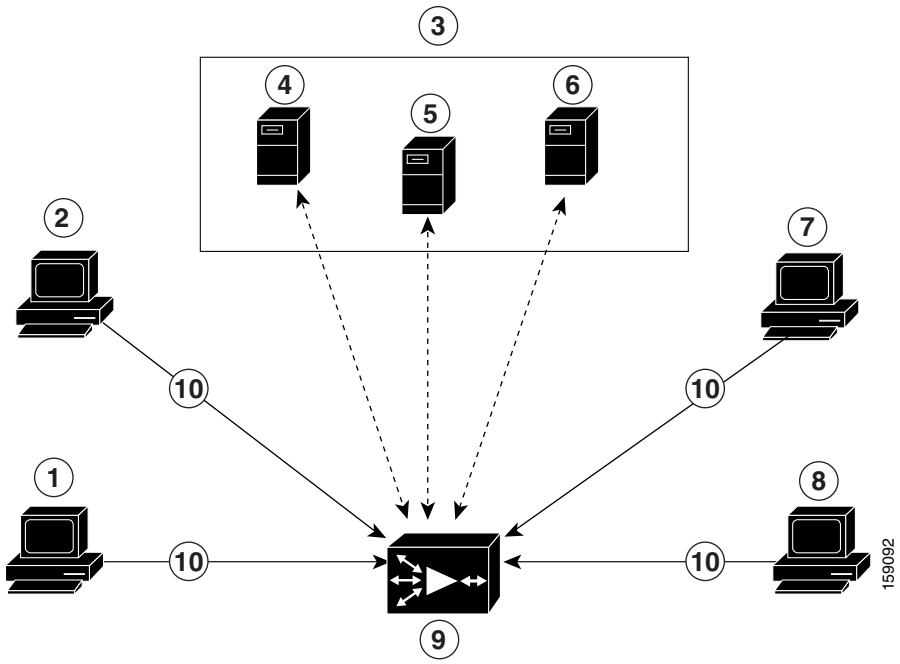

1	FTP/SFTP client	**6**	Windows domain server
2	WAAS Central Manager GUI or WAE Device Manager GUI	**7**	Console or Telnet clients
3	Third-party AAA servers	**8**	SSH client
4	RADIUS server	**9**	WAE that contains a local database and the default primary authentication database
5	TACACS+ server	**10**	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 user account in the WAAS Central Manager as described in Chapter 7, "Creating and Managing Administrator User Accounts and Groups." The user account must not be a local one; that is, do not check the Local User check box when you create the account.

For more information on the default AAA configuration, see the "Default Administrative Login Authentication and Authorization Configuration" section on page 6-4. For more information on configuring AAA, see the "Configuring Administrative Login Authentication and Authorization" section on page 6-6.

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 6-1 lists the default configuration for administrative login authentication and authorization.

Table 6-1 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+ 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

Table 6-1 Default Configuration for Administrative Login Authentication and Authorization (continued)

Feature	Default Value
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
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 6-6.

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.

To invoke these utilities from the WAAS Central Manager GUI, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the device for which you want to run the utilities in a predefined order.

Step 3 In the navigation pane, choose **Configure > Security > AAA > Windows Domain**.

Step 4 In the displayed window, click the **Show Authentication Status** button at the bottom of the window.

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 6-7.

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 6-13
- About TACACS+ Server Authentication Settings, page 6-15
- Configuring Windows Domain Server Authentication Settings, page 6-17

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 6-26.

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 6-1). 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, page 6-7
- Configuring Remote Authentication Server Settings for WAAS Devices, page 6-13
- Enabling Administrative Login Authentication and Authorization Schemes for WAAS Devices, page 6-26

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, page 6-8
- Disabling and Reenabling the Telnet Service for WAAS Devices, page 6-10
- Configuring Message of the Day Settings for WAAS Devices, page 6-10
- Configuring Exec Timeout Settings for WAAS Devices, page 6-11
- Configuring Line Console Carrier Detection for WAAS Devices, page 6-12

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device (or device group) for which you want to enable SSH.

Step 3 In the navigation pane, choose **Configure > Network > Console Access > SSH**.

The SSH Configuration window appears. (See Figure 6-3.)

Figure 6-3 SSH Configuration Window

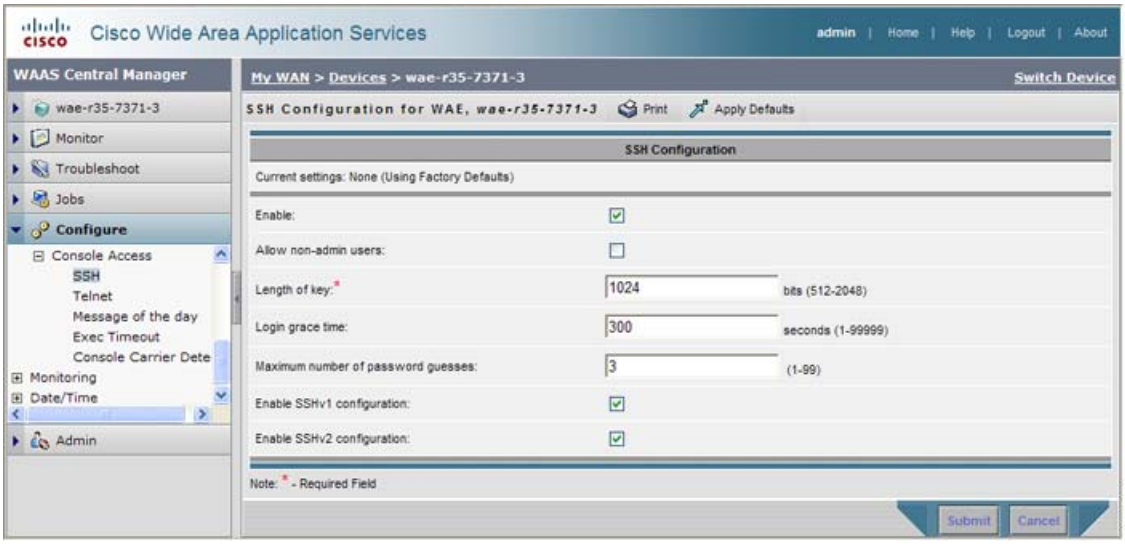

Step 4 Check the **Enable** check box to enable the SSH feature. SSH enables login access to the chosen WAAS device (or the device group) through a secure and encrypted channel.

Step 5 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 6 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 7 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 8 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 9 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 4.)

Step 10 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 7 or 8. If you want to use the Telnet button from the Device Dashboard, use Internet Explorer 6 or 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device (or device group) for which you want to disable Telnet.

Step 3 In the navigation pane, choose **Configure > Network > Console Access > Telnet**. The Telnet Settings window appears.

Step 4 Uncheck the **Telnet Enable** check box to disable the terminal emulation protocol for remote terminal connection for the chosen device (or device group).

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 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 GUI navigation pane, choose **My WAN > Manage Devices**. The Devices window appears.

Step 2 Click the **Edit** icon next to the WAAS device for which you want to set up the message of the day. The Device Dashboard window for the chosen device appears.

Step 3 In the navigation pane, choose **Configure > Network > Console Access > Message of the day**. The MOTD Configuration window for the chosen device appears.

Step 4 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 5 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 6 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 7 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 8 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device (or device group) for which you want to configure the EXEC timeout.

Step 3 In the navigation pane, 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device (or device group) that you want to configure.

Step 3 In the navigation pane, choose **Configure > Network > Console Access > Console Carrier Detect**. The Console Carrier Detect Settings window appears.

Step 4 Check the **Enable console line carrier detection before writing to the console** check box to enable the window for configuration.

Step 5 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 6 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 6-13
- About TACACS+ Server Authentication Settings, page 6-15
- Configuring TACACS+ Server Settings, page 6-16
- Configuring Windows Domain Server Authentication Settings, page 6-17
- LDAP Server Signing, page 6-24

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or device group) that you want to configure.

Step 3 From the navigation pane, choose **Configure > Security > AAA > RADIUS**. The RADIUS Server Settings window appears. (See Figure 6-4.)

Figure 6-4 RADIUS Server Settings Window

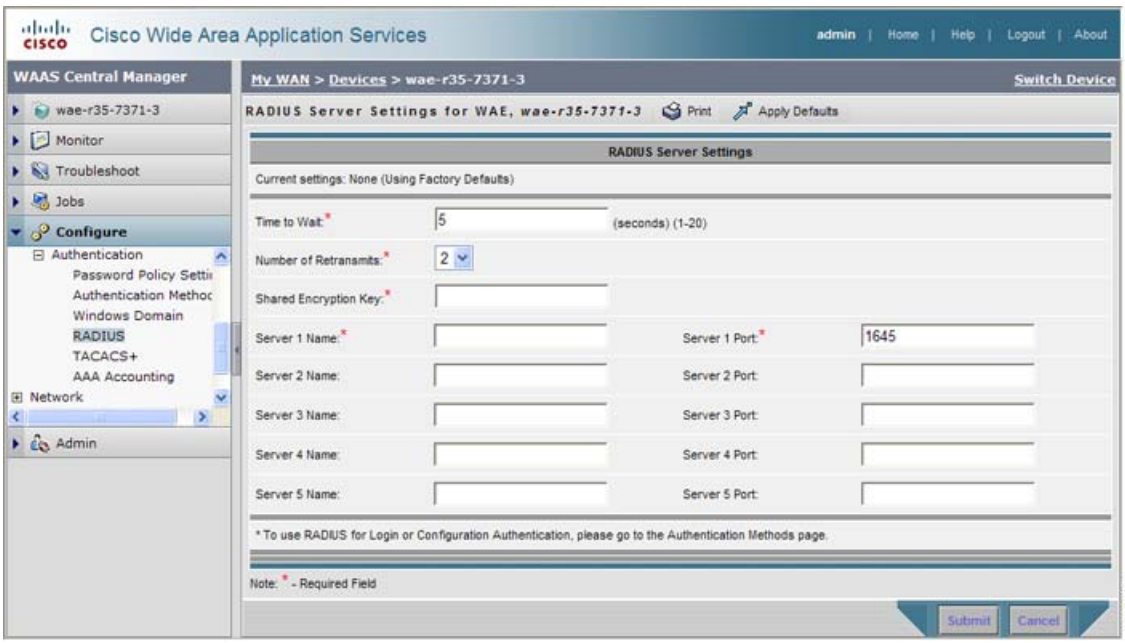

Step 4 In the Time to Wait field, specify how long the device or device group should wait for a response from the RADIUS server before timing out. The range is from 1 to 20 seconds. The default value is 5 seconds.

Step 5 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 6 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: backwards single quote (`), double quote ("), pipe (|), closing bracket (]), or number sign (#).

Step 7 In the Server Name field, enter an IP address or hostname of the RADIUS server. Five different hosts are allowed.

Step 8 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 9 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 6-26.

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 7-3.) 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 but not in 5.0.

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.

Note When using TACACS+ user authentication, usernames cannot begin with numbers or consist of only numbers, otherwise login will fail.

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 a privileged EXEC password (superuser or admin-equivalent 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or device group) that you want to configure.

Step 3 From the navigation pane, 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 4 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 5 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 6 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 7 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: backwards single quote (`), double quote ("), pipe (|), closing bracket (]), or number sign (#).

Step 8 In the Primary Server field, enter an IP address or hostname for the primary TACACS+ server.

Step 9 In the Secondary Server field, enter an IP address or hostname for a secondary TACACS+ server.

Step 10 In the Tertiary Server field, enter an IP address or hostname for a tertiary TACACS+ server.

Note You can specify up to two backup TACACS+ servers.

Step 11 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 6-26.

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
- Disconnected mode operation

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 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.")

This section contains the following topics:

- Configuring Windows Domain Server Settings on a WAAS Device, page 6-18
- Unregistering a WAE from a Windows Domain Controller, page 6-22
- Disabling the Automatic Machine Account Password Changes for the Edge WAE, page 6-23

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.

To configure Windows Domain server settings on a WAAS device or device group, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or device group) that you want to configure.

Step 3 From the navigation pane, choose **Configure > Security > AAA > Windows Domain**. The Windows Domain Server Settings window appears. (See Figure 6-5.)

Figure 6-5 Windows Domain Server Settings Window

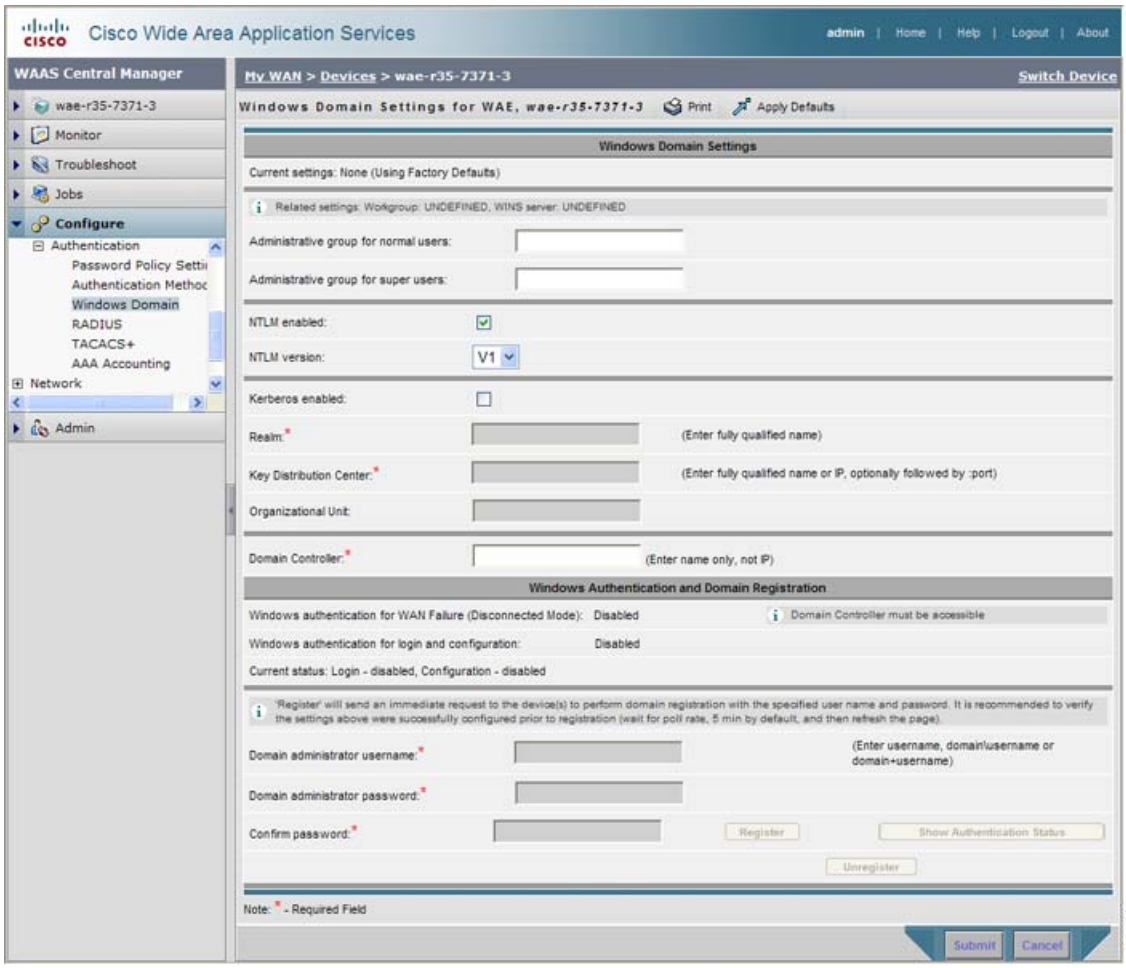

Note If the related WINS server and the workgroup or domain name have not been defined for the chosen device (or device group), an informational message is displayed at the top of this window to inform you that these related settings are currently not defined, as shown in Figure 6-5. To define these settings, choose **Configure > Network > WINS**.

Step 4 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 only has a privilege level of 0, by entering the name of the group in the Administrative group for normal users field.

Note By default, there are not predefined user groups for Windows domain authorization configured on a WAE.

Step 5 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, by entering the name of the group in the Administrative group for superusers field.

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 6 Select NTLM or Kerberos as a shared secure authentication method for administrative logins to the chosen device (or device group) as follows:

Note Kerberos version 5 is used for Windows systems running Windows 2000 or higher with users logging in to domain accounts.

- To enable NTLM, check the **NTLM enabled** check box.
- To select NTLM version 1, check the **NTLM enabled** check box. NTLM version 1 is selected by default.

  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.
- To select NTLM version 2, choose **V2** from the drop-down list.

  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.

- To select Kerberos, check the **Kerberos enabled** check box. In the 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 desired, enter the name of the organizational unit in the Organizational Unit field.

  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 7 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 8 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 9, are not stored in the WAAS Central Manager's database.

Step 9 Register the chosen device (or device group) with the Windows Domain Controller as follows:

a. In the Domain Administrator username field, enter the administrative username (the domain\username or the domain name plus the username) of the specified Windows Domain Controller.

b. In the Domain Administrator password field, enter the administrative password of the specified Windows Domain Controller.

c. In the Confirm password field, reenter the administrative password of the specified Windows Domain Controller.

d. Click the **Register** button.

Note When you click the Register button, the WAAS Central Manager immediately sends a registration request to the WAAS device (or 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 administrator's 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).

e. To check the status of the registration request, you must set Windows as the authentication and authorization method for the device, as described in the "Enabling Administrative Login Authentication and Authorization Schemes for WAAS Devices" section on page 6-26, then wait a few minutes and click the **Show Authentication Status** button.

You may choose to check the **Refresh Authentication Status** check box before you click the **Show Authentication Status** button. When this box is checked, the WAAS Central Manager queries the device for an updated domain registration status. When this box is unchecked, the status is retrieved from the Central Manager local cache.

After you click the **Show Authentication Status** button, a dialog box appears prompting if you want to continue with this request to view the status of the authentication request. (See Figure 6-6.)

Figure 6-6 Confirmation Dialog Box

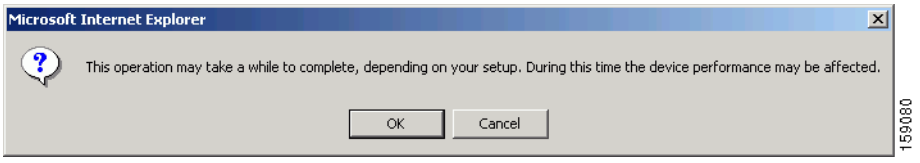

f. Click **OK** to continue or click **Cancel** to cancel the request.

If the request fails, you will receive an error dialog. Wait a few more minutes and try again to see the updated authentication status.

If the request succeeds, the domain registration status will be shown immediately below the Windows Authentication and Domain Registration heading, in the lower part of Figure 6-5. In addition, the status of windows authentication and disconnected mode is shown in this area.

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 6-26.

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** 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, and then verify it, by using the following commands (for Kerberos authentication):

```
WAE# windows-domain diagnostics net  "ads join -U AdminUsername%AdminPassword"
WAE# windows-domain diagnostics net  "ads testjoin -U AdminUsername%AdminPassword"
```

For NTLM authentication, use the following commands instead:

```
WAE# windows-domain diagnostics net  "rpc join -U AdminUsername%AdminPassword"
WAE# windows-domain diagnostics net  "rpc testjoin -U AdminUsername%AdminPassword"
```

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
```

To enable content request authentication in disconnected mode, use the **authentication content-request windows-domain disconnected-mode enable** configuration command.

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or the device group) that you want to unregister.

Step 3 From the navigation pane, choose **Configure > Security > AAA > Authentication Methods**. The Authentication and Authorization Methods window appears. (See Figure 6-7 on page 6-29.)

Step 4 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 6-26.

Step 5 Click **Submit** to save the settings.

Step 6 From the navigation pane, choose **Configure > Security > AAA > Windows Domain**. The Windows Domain Server Settings window appears. (See Figure 6-5.)

Step 7 (Optional) Enter the administrative username and password in the Domain administrator username, Domain administrator 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 8 Scroll down and click the **Unregister** button.

Note When you click the Unregister 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.

a. Check the status of the unregistration request by waiting a few minutes and clicking the **Show Authentication Status** button. A dialog box appears prompting if you want to continue with this request to view the status of the authentication request. (See Figure 6-6.)

b. Click **OK** to continue or click **Cancel** to cancel the request.

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 diagnostics net  "ads leave -U AdminUsername%AdminPassword"
```

There is no CLI command to unregister the WAAS device if it is using NTLM authentication.

Disabling the Automatic Machine Account Password Changes for the Edge WAE

In a WAAS network where a Windows domain controller is configured for authentication and Disconnected Mode is enabled on an Edge WAE operating in WAFS legacy mode, the domain controller authenticates content requests in the event of a WAN failure. By default, Windows domain controllers enforce automatic machine account password changes as part of the authentication process. The machine account password for the Edge WAE is automatically negotiated and changed between the Edge WAE and the domain controller every seven days. However, if the authentication service is down, this process may not occur, and the machine account password for the Edge WAE may expire.

To prevent this situation, we recommend that you disable automatic machine account password changes for the Edge WAE. The procedure that follows describes how to disable automatic machine account password changes for Windows XP and Windows Server 2003 using Group Policy Editor. Refer to Microsoft's Help and Support page for details on how to disable automatic machine account password changes for the other Windows operating systems.

To disable the automatic machine account password changes for the Edge WAE using Group Policy Editor, follow these steps:

Step 1 On the Windows domain controller, click **Start** and then choose **Run**.

Step 2 Enter **Gpedit** at the prompt, and then click **OK**.

Step 3 Expand the Local Computer Policy, Windows Settings, Security Settings, Local Policies, Security Settings, Local Policies, Security Options.

Step 4 Configure the following setting: Domain Member: Disable machine account password changes (DisablePasswordChange).

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 6-25
- Disabling LDAP Server Signing on a Client WAE, page 6-26

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 terminal
WAE(config)# smb-conf section "global" name "ldap ssl" value "start_tls"
```

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 Administrator%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 "start_tls"
```

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 6-4.

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 6-15, the "Configuring RADIUS Server Authentication Settings" section on page 6-13, and the "Configuring Windows Domain Server Authentication Settings" section on page 6-17 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 **My WAN > Manage Devices.** Click the **Edit** icon next to the name of the WAAS device for which you want to change the default login authentication failover method, and then choose **Configure > Security > AAA > Authentication Methods** from the navigation pane. 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 **My WAN > Manage Device Groups.** Click the **Edit** icon next to the name of the device group for which you want to change the default login authentication failover method, and then choose **Configure > Security > AAA > Authentication Methods** from the navigation pane. 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.

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or the device group) that you want to configure.

Step 3 From the navigation pane, choose **Configure > Security > AAA > Authentication Methods**. The Authentication and Authorization Methods window appears. (See Figure 6-7.)

Figure 6-7 Authentication and Authorization Methods Window

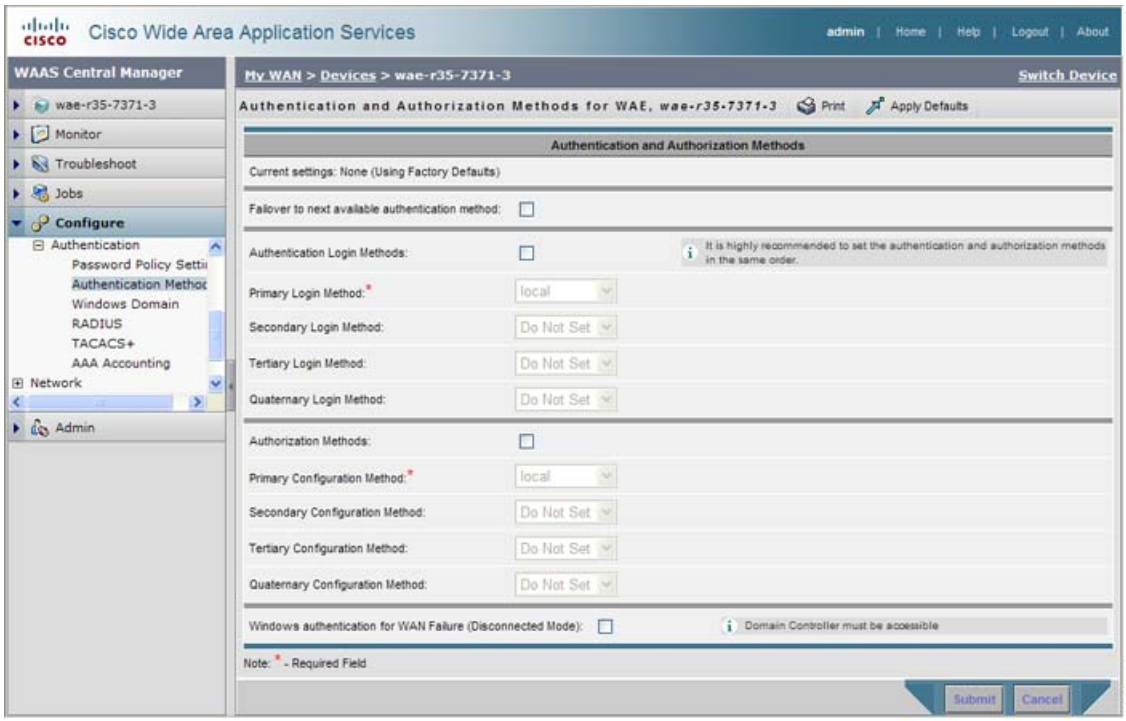

Step 4 Check the **Failover to next available authentication method** check box to query the secondary authentication database only if the primary authentication server is unreachable. When the box is unchecked, the other authentication methods are tried if the primary method fails for any reason.

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 5 Check the **Authentication Login Methods** check box to enable authentication privileges using the local, TACACS+, RADIUS, or WINDOWS databases.

Step 6 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 7 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 8 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 4), 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 4), 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 9 Check the **Windows authentication for WAN Failure (Disconnected Mode)** check box to enable content request authentication in disconnected mode. Disconnected mode is available only if you are using WAFS legacy mode.

When this feature is enabled, the Windows domain server authenticates the content request in disconnected mode. By default, this feature is disabled on a WAE.

If you enable this feature, we recommend that you also disable automatic account password changes for the WAE. For more information, see the "Disabling the Automatic Machine Account Password Changes for the Edge WAE" section on page 6-23.

Note Windows authentication for disconnected mode operates only if you select NTLM as the shared secure authentication method in the Windows Domain settings window, as described in the "Configuring Windows Domain Server Authentication Settings" section on page 6-17.

Step 10 Click **Submit** to save the settings.

Note If you have enabled the Windows authentication or authorization method, the Central Manager queries the WAE 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 device group, the Central Manager does not query the individual WAEs 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or the device group) that you want to configure.

Step 3 From the navigation pane, choose **Configure > Security > AAA > Command Authorization**. The Command Authorization window appears. (See Figure 6-8.)

Figure 6-8 Command Authorization Settings Window

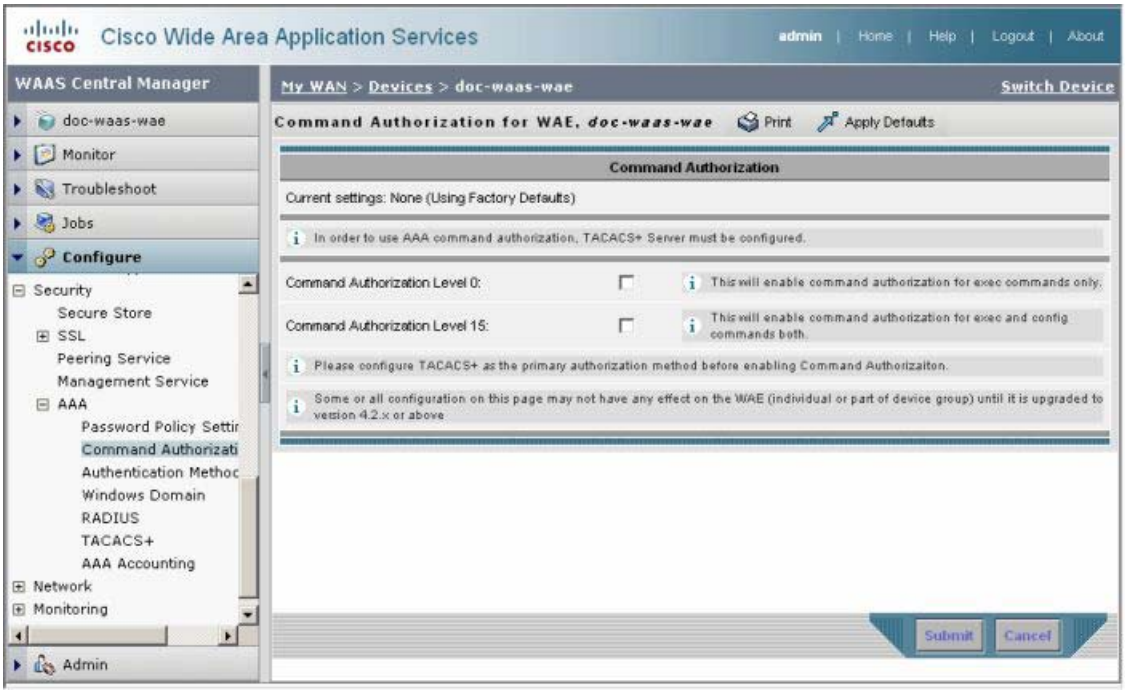

Step 4 Check the Command Authorization Level check box for the desired level.

- 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 5 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 6-15.)

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 8, "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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or the device group) that you want to configure.

Step 3 In the navigation pane, choose **Configure > Security > AAA > AAA Accounting**. The AAA Accounting Settings window appears.

Step 4 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 5 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 6 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 7 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 6-2 describes the event type options.

Table 6-2 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.
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 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 16-62.

CHAPTER 7

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, page 7-1
- Creating and Managing User Accounts, page 7-2

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 7-1 for a description of the two types of accounts you can create from the WAAS Central Manager GUI.

Table 7-1 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 drawer in the WAAS Central Manager GUI.
Local account	Provides CLI access to WAE devices and optionally allows users to access the Print Services Administration GUI and 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, page 7-2
- Working with Accounts, page 7-3
- Working with Passwords, page 7-9
- Working with Roles, page 7-10
- Working with Domains, page 7-14
- Working with User Groups, page 7-18

Overview for Creating an Account

Table 7-2 provides an overview of the steps you must complete to create a new roles-based administrator account.

Table 7-2 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 7-4.

Table 7-2 Checklist for Creating a Roles-based Administrator Account (continued)

Task	Additional Information and Instructions
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 7-11. If you are using an external authentication server, you can define matching user groups that automatically assign roles to users.
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 7-13. 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 or device groups that the new account can manage. For more information, see the "Creating a New Domain" section on page 7-15.
5. Add an entity to the domain.	Adds one or more WAEs or device groups to the domain. For more information, see the "Adding an Entity to a Domain" section on page 7-15.
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 7-16. 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 two predefined roles, known as the admin and print roles. The admin role has access to all services. The print role has access to all print related pages. 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.

Note When using external authentication of users on a TACACS+ server, usernames may not begin with numbers or consist of only numbers, otherwise login will fail.

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:

- Creating a New Account, page 7-4
- Modifying and Deleting User Accounts, page 7-6
- Changing the Password for Your Own Account, page 7-7
- Changing the Password for Another Account, page 7-8
- Viewing User Accounts, page 7-8
- Unlocking User Accounts, page 7-9

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 7-3 describes the results of creating a local CLI user when setting up an account.

Table 7-3 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 GUI navigation pane, 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 7-3 on page 7-4 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 Check the **Print Admin** check box to use this account to upload drivers to the central repository on the WAAS Central Manager and to access the Print Services Administration GUI.

For more information, see the "Setting Up the WAAS Central Manager as the Driver Repository" section on page 13-16 and the "Using the Print Services Administration GUI" section on page 13-28.

Note the following about the print admin account:

- This Print Admin check box is enabled only after you check the **Local User** check box.
- The print admin account must have a privilege level of 15 (super user) in order to use the account to upload drivers to the repository. If the print admin account has a privilege level of 0, it can be used only to access the Print Services Administration GUI.
- The print admin account does not have access to print related pages in the WAAS Central Manager unless it also has the print or admin roles assigned.
- You must assign a domain to the print admin account. All WAEs to which the user needs access should be members in the assigned domain.
- You cannot assign the print role to an externally authenticated user.

Step 7 (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 8 (Optional) In the Comments field. enter any additional information about this account.

Step 9 Click **Submit**.

A Changes Submitted message appears at the bottom of the window.

Step 10 Assign roles to this new account as described in the "Working with Roles" section on page 7-10 and assign domains as described in the "Working with Domains" section on page 7-14.

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 GUI navigation pane, 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 GUI navigation pane, choose **Admin > 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 GUI navigation pane, 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 7-4.

To view user accounts for a specific device, choose **My WAN > Manage Devices** and click the **Edit** icon next to the device for which you want to view users. Then click **Device Users** or **CM Users**, depending on whether the device is an application accelerator or a Central Manager. 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group on which you want to change the password policy.

Step 3 From the WAAS Central Manager GUI navigation pane, choose **Configure > Security > AAA > Password Policy Settings**.

Step 4 Check the **Enforce stringent password** check box to enable the strong password policy.

Step 5 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 7-9.

Step 6 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 7-9.

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 two predefined roles, known as the admin and print roles. 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. The print role has access to all print related pages in the WAAS Central Manager.

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 7-18.

This section contains the following topics:

- Creating a New Role, page 7-11
- Assigning a Role to a User Account, page 7-13
- Modifying and Deleting Roles, page 7-13
- Viewing Role Settings, page 7-14

Creating a New Role

To create a new role, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, 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 7-4 lists the services that you can enable for a role.

Table 7-4 Description of the WAAS Services

Service	Description
Devices	Allows this role to configure and manage the settings in the My WAN drawer of the WAAS Central Manager GUI. If you do not want to enable the entire My WAN drawer, select the subpages that you want this role to manage.
Monitor	Allows this role to configure and manage the settings in the Monitor drawer of the WAAS Central Manager GUI. If you do not want to enable the entire Monitor drawer, select the subpages that you want this role to manage.

Table 7-4 Description of the WAAS Services (continued)

Service	Description
Report	Allows this role to configure and manage the settings in the Report drawer of the WAAS Central Manager GUI. If you do not want to enable the entire Report drawer, select the subpages that you want this role to manage.
Jobs	Allows this role to configure and manage the settings in the Jobs drawer of the WAAS Central Manager GUI. If you do not want to enable the entire Jobs drawer, select the subpages that you want this role to manage.
Configure	Allows this role to configure and manage the settings in the Configure drawer of the WAAS Central Manager GUI. If you do not want to enable the entire Configure drawer, select the subpages that you want this role to manage.
Admin	Allows this role to access the items in the Admin drawer of the WAAS Central Manager GUI. If you do not want to enable the entire Admin drawer, select the subpages that you want this role to manage.
All WAEs	Allows this role to access all the WAEs in your WAAS network. If this service is not enabled, the user account will only have access to the WAEs 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 WAEs in your network. • Assigning to the account the domain that contains all the WAEs.
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.
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 window, My WAN > Alerts. For more information about device alarms, see Chapter 16, "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.

Note If you create a role that has access to the Configure > File Services > Baseline Group, Configure > Acceleration > Baseline Group, or Configure > Platform > Baseline Group pages, this role must also have the All Device Groups service enabled, in order to be able to access all necessary configuration pages. Alternatively, a user with this role could also be assigned a domain that includes all devices that are part of the baseline groups.

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 GUI navigation pane, 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 GUI navigation pane, 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 GUI navigation pane, 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 *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 by the user.

When you create a domain, you choose the type of entities that can be associated with the domain. Entity types include WAEs, 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" section on page 11-21.

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 7-18.

This section contains the following topics:

- Creating a New Domain, page 7-15
- Adding an Entity to a Domain, page 7-15
- Assigning a Domain to a User Account, page 7-16
- Modifying and Deleting Domains, page 7-17
- Viewing Domains, page 7-17

Creating a New Domain

To create a new domain, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, 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 WAEs, 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" section on page 11-21.

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 WAEs 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 GUI navigation pane, 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 group specifies the entities (devices or device groups) or file server dynamic shares that the account or group can manage.

Note If the role that you assigned to an account or group has the All WAEs 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 WAEs and/or device groups in the WAAS system. For more information, see Table 7-4 on page 7-11.

To assign a domain to a user account or group, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, 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 GUI navigation pane, 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 GUI navigation pane, 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 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; instead, you define roles and domains for the user groups, and the user is assigned the roles and domains that are defined for the groups to which they belong.

Note 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 but not in 5.0.

WAAS reads group membership information for each user from the external authentication server.

This section contains the following topics:

- Creating a New User Group, page 7-18
- Assigning Roles to a User Group, page 7-19
- Assigning Domains to a User Group, page 7-19
- Modifying and Deleting a User Group, page 7-20
- Viewing User Groups, page 7-20

Creating a New User Group

To create a new user group, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, 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 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 GUI navigation pane, 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 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 WAEs 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 WAEs and/or device groups in the WAAS system. For more information, see Table 7-4 on page 7-11.

To assign a domain to a user group, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, 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 GUI navigation pane, 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 7-18.

CHAPTER 8

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:

- About IP ACLs for WAAS Devices, page 8-1
- Creating and Managing IP ACLs for WAAS Devices, page 8-2
- List of Extended IP ACL Conditions, page 8-7

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, and standby interfaces, and on Cisco WAE Inline Network Adapter InlineGroup interfaces. This type of ACL is intended to control management traffic (telnet, SSH, and Central Manager GUI). Use the **ip access-group** interface configuration command to apply an interface ACL.
- Interception ACL—Applied on all interfaces. 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 an Interception Access Control List" section on page 4-29.

- 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 8-2.

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 GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the name of the device (for example, a data center WAE named bd-s14) for which you want to create an IP ACL.

Step 3 In the navigation pane, 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 4 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 5 Click **Submit** to save the IP ACL named test1. IP ACLs without any conditions defined do not appear on the individual devices.

Step 6 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 8-1.)

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 8-1 Creating a New Condition for an Extended IP ACL Window

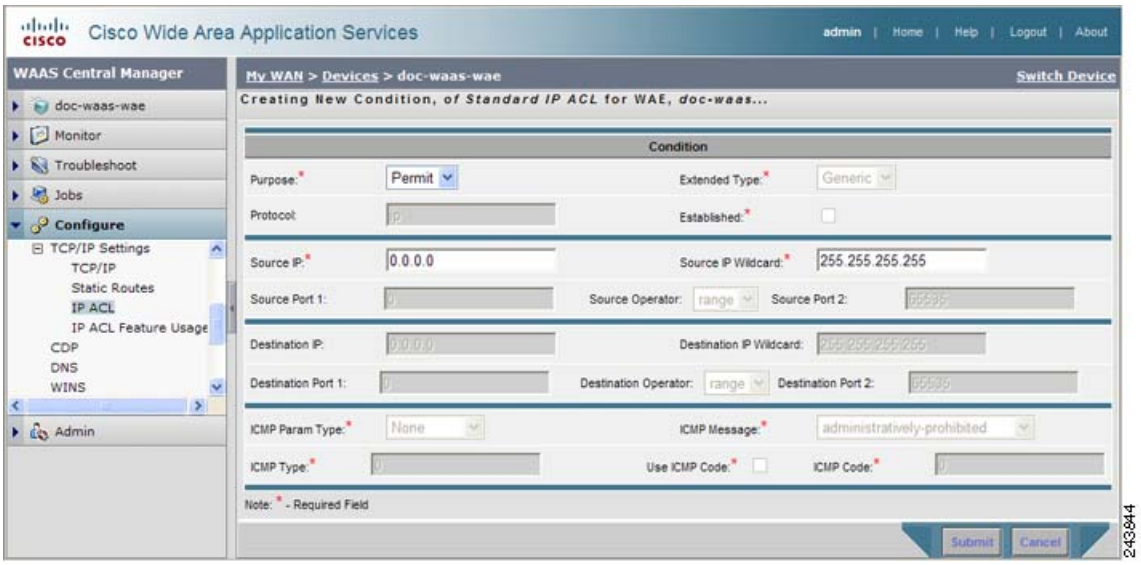

b. Enter values for the properties that are enabled for the type of IP ACL that you are creating, as follows:

- To set up conditions for a standard IP ACL, go to Step 7.
- To set up conditions for an extended IP ACL, go to Step 8.

Step 7 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 8-1 describes the fields in a standard IP ACL.

Table 8-1 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*	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*	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 8 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 8-2.)

Table 8-2 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*	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 8-4 through Table 8-7.)

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 9 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 10 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. In the navigation pane, 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 8-3.) If you do not want to associate an IP ACL with one of the applications, choose **Do Not Set**.

Table 8-3 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 11 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. In the navigation pane, choose **Configure > Network > Network Interfaces**.

The Network Interfaces window for the device appears. This window displays all the interfaces available on that device.

Note The Port Type column may contain a port-channel interface indicating an EtherChannel configuration. EtherChannel for the WAAS software supports the grouping of up to four same-speed network interfaces into one virtual interface.

c. Click the **Edit** icon next to the name of the interface to which you want to apply an IP ACL. The Modifying Network Interface window appears.

d. Scroll to the bottom of the window. From the Inbound ACL drop-down list, 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 12 Click **Submit** to save the settings.

Step 13 To use an IP ACL to define the traffic that should be intercepted, see the "Configuring an Interception Access Control List" section on page 4-29.

Step 14 (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. In the navigation pane, 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 8 in the "Creating and Managing IP ACLs for WAAS Devices" section on page 8-2).

The list of extended IP ACL conditions are as follows:

- Generic (See Table 8-4.)
- TCP (See Table 8-5.)

- UDP (See Table 8-6.)
- ICMP (See Table 8-7.)

Table 8-4 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*	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*	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*	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 8-5 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*	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*	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*	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 8-5 Extended IP ACL TCP Condition (continued)

Field	Default Value	Description
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.
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 8-6 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*	UDP	Matches the UDP Internet protocol.
Established	—	Not available for UDP.
Source IP*	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*	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 8-6 Extended IP ACL UDP Condition (continued)

Field	Default Value	Description
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.
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 8-7 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*	ICMP	Matches the ICMP Internet protocol.
Source IP*	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*	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.

Table 8-7 Extended IP ACL ICMP Condition (continued)

Field	Default Value	Description
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.
ICMP Param Type*	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*	administratively-prohibited	Allows a combination of ICMP type and code to be specified using a keyword chosen from the drop-down list.
ICMP Type*	0	Number from 0 to 255. This field is enabled when you choose **Type/Code**.
Use ICMP Code*	Unchecked	When checked, enables the ICMP Code field.
ICMP Code*	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 9

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.

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, page 9-1
- Managing Software Licenses, page 9-3
- Enabling the Inetd RCP and FTP Services, page 9-4
- Configuring Date and Time Settings, page 9-5
- Configuring Secure Store Settings, page 9-10
- Modifying the Default System Configuration Properties, page 9-17
- Configuring Web Application Filter, page 9-19
- Configuring Faster Detection of Offline WAAS Devices, page 9-22
- Configuring Alarm Overload Detection, page 9-23
- Configuring the E-mail Notification Server, page 9-24

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 GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the device that you want to modify.

The Device Dashboard window appears.

Step 3 In the navigation pane, 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 4 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.

Note If the WAE is operating in nontransparent mode and print services is enabled, you must configure identical names for the NetBIOS name and the hostname of the device that you enter in the Name field.

Step 5 Under the Locality heading, set or change the location by choosing a new location from the **Location** drop-down list. To create a new location for this device, see the "Creating Locations" section on page 3-14.

Step 6 Under the NAT Configuration 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.
- 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.

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 7 In the Comments field, enter any comments that you want to appear for this device.

Step 8 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 9-1 lists the software licenses that may be purchased and the features that each license enables.

Table 9-1 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 (WAFS), 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.

To add a license to a WAE from the WAAS Central Manager, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the WAE device that you want to modify. (Do not choose a Central Manager device because you must use the CLI to manage licenses on Central Managers.)

Step 3 In the navigation pane, choose **Admin > License Management**.

Step 4 Check the check box next to each license that you want to add.

Step 5 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group for which you want to enable RCP services.

Step 3 In the navigation pane, choose **Configure > Network > Network Services**. The Network Services window appears.

Step 4 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 5 Check the **Enable FTP Service** check box to enable the Inetd FTP service. By default, this option is disabled.

Step 6 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, page 9-5
- Configuring Time Zone Settings, page 9-5

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group that you want to configure.

Step 3 From the navigation pane, choose **Configure > Date/Time > NTP**. The NTP Settings window appears.

Step 4 Check the **Enable** check box to enable NTP settings. By default, this option is disabled.

Step 5 In the NTP Server field, enter a hostname or IP address.

Step 6 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group for which you want to configure the time zone.

Step 3 In the navigation pane, choose **Configure > Date/Time > Time Zone**. The Time Zone Settings window appears.

Step 4 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 9-2.)

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 9-3.

Step 5 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 6 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 7 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 8 In the Offset field, specify the minutes offset from UTC (0–1439). (See Table 9-3.)

The summer time offset specifies that the number of minutes that the system clock moves forward at the specified start time and backward at the end time.

Step 9 Click the **No Customized Summer Time Configured** radio button to not specify a summer or daylight saving time for the corresponding time zone.

Step 10 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 9-2 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 9-2 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 9-3 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 9-3 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 9-3 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 stronger 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, page 9-10
- Enabling Secure Store Encryption on the Central Manager, page 9-12
- Enabling Secure Store Encryption on a Standby Central Manager, page 9-13
- Enabling Secure Store Encryption on a WAE Device, page 9-13
- Changing the Secure Store Encryption Key and Password, page 9-14
- Resetting Secure Store Encryption on a Central Manager, page 9-15
- Disabling Secure Store Encryption on a WAE Device, page 9-16

Secure Store Overview

When you enable 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 passwords, user login passwords, and certificate key files.

To enable secure store encryption, you must enter a password on the Central Manager. 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, the WAFS core password, and user passwords.

When secure store is enabled on the Central Manager, the data is encrypted using a 256-bit key encryption key generated from the password you enter 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.

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, you must re-enable secure store encryption manually. Disk encryption and CIFS preposition features on the remote WAE devices will not operate properly until you enter the secure store password on the Central Manager to re-enable secure store encryption.

When secure store is enabled, the following system characteristics are affected:

- 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 will be handled by the standby Central Manager. (Secure store mode must be enabled manually on the standby Central Manager.)
- Backup scripts will backup the secure store mode status of the device at the time of backup. Backup is supported only on the Central Manager.
- Restore scripts will verify if the backup file has secure store mode enabled. If the backup file is in secure store mode, you must enter the pass phrase to verify and restore the device. Restore is supported only on the Central Manager.
- 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 will not include preposition configuration in the updates that it sends to the Central Manager.

- If secure store is active, you cannot downgrade to an earlier version of WAAS software that does not support secure store mode. You must disable secure store mode before installing the previous version of WAAS software.
- 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 Enabling Disk Encryption, page 15-28.

Enabling Secure Store Encryption on the Central Manager

To enable secure store encryption on the Central Manager, follow these steps:

Step 1 From the WAAS Central Manager GUI, choose **Admin > Secure Store**. The Configure CM Secure Store window appears.

Step 2 Enter a password in the Enter passphrase and Confirm passphrase fields.

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 **Initialize** button.

The secure store is initialized and opened. Data is encrypted using the key derived from the password.

To enable secure store from the CLI, use the **cms secure-store init** EXEC command.

Note Whenever you reboot the Central Manager, you must reopen the secure store manually. Disk encryption and CIFS preposition features on the remote WAE devices will not operate properly until you enter the secure store password on the Central Manager to reopen the secure store. Use the same configuration screen as above, except you should use the Open Secure Store section. The secure store is already initialized, so the Initialize Secure Store section is not shown. Alternatively, you can use the **cms secure-store open** EXEC command.

Note When you enable secure store on the primary Central Manager, you should enable secure store on the standby Central Manager as well. See Enabling Secure Store Encryption on a Standby Central Manager, page 9-13.

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.

To enable secure store encryption on a standby Central Manager, first enable 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 Enabling Secure Store Encryption on the Central Manager, page 9-12.

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 GUI, choose **Manage Devices** (or **Manage Device Groups**).

Step 2 Click the Edit icon next to the device or device group for which you want to enable secure store.

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 Working with Device Groups, page 3-2.

Step 3 From the navigation pane, choose **Configure > Security > Secure Store**. The Secure Store Settings window appears, as shown in Figure 9-1.

Figure 9-1 Example of Secure Store Settings Window

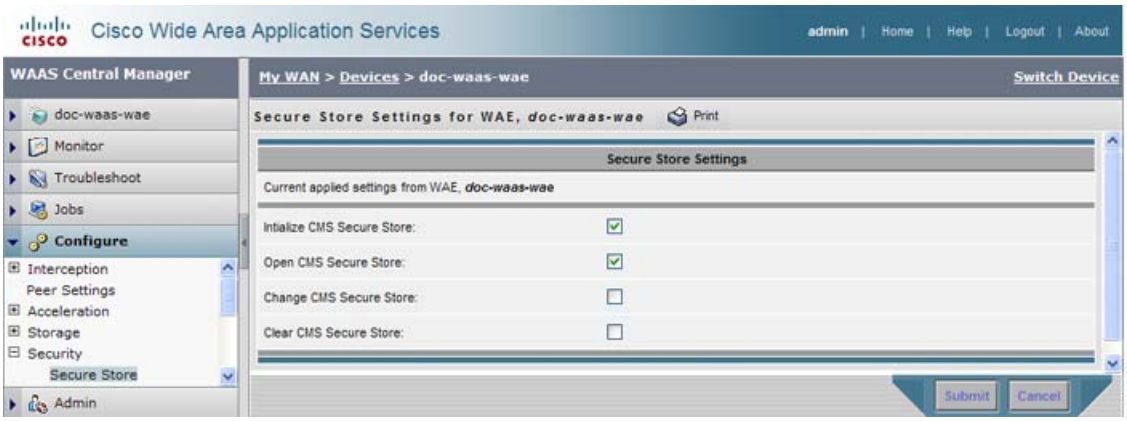

Step 4 Check the **Initialize CMS Secure Store** box. (The Open CMS Secure Store box will be checked automatically.)

Step 5 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 will be lost and you will need to 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 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.

To change the password and generate a new encryption key on the Central Manager, follow these steps:

Step 1 From the WAAS Central Manager GUI, choose **Admin > Secure Store**.

Step 2 In the Current passphrase field, enter the current password.

Step 3 In the Enter 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 use the WAAS Central Manager GUI and do the following:

Step 1 From the WAAS Central Manager GUI, choose **Manage Devices** (or **Manage Device Groups**).

Step 2 Click the Edit icon next to the device or device group for which you want to generate a new encryption key.

Step 3 From the navigation pane, choose **Configure > Security > Secure Store**.

Step 4 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

Use the **cms secure-store reset** command if you reload the Central Manager and you cannot open the secure store because you forget the secure store password. This command deletes all encrypted data, certificate and key files, and key manager keys. The secure store is left in the uninitialized state.

To reset secure store encryption on a Central Manager, follow these steps:

Step 1 Enter the **cms secure-store reset** command on the primary Central Manager.

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/or CIFS legacy mode core passwords.

For information on resetting user passwords, see the "Changing the Password for Another Account" section on page 7-8. For information on resetting CIFS legacy mode core cluster passwords, see the "Configuring the Core Cluster" section on page 11-11. For information on resetting dynamic share passwords, see the "Creating Dynamic Shares" section on page 11-21. For information on resetting preposition passwords, see the "Creating a Preposition Directive" section on page 11-25.

Step 5 Reinitialize and open secure store on the primary Central Manager as described in the "Enabling Secure Store Encryption on the Central Manager" section on page 9-12.

Step 6 If secure store on the standby Central Manager is initialized but not open, wait until the primary Central Manager replicates data to the standby Central Manager, then open secure store on the standby Central Manager using the **cms secure-store open** command.

Step 7 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 9-16). 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 9-13). Or, from the CLI, enter the **cms secure-store init** EXEC command. (This step is needed only if you performed step 7a.)

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 15-28). Or, from the CLI, enter the **no disk encrypt enable** global configuration command.

e. If disk encryption had been enabled before step 7d, reload the device. After the reload, reenable disk encryption and reload the device again.

Note If the WAE is reloaded before doing Step 7, disk encryption, SSL acceleration, and secure store will not function properly. In this case, you must restore the WAE to factory defaults.

Step 8 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 GUI, choose **Manage Devices** (or **Manage Device Groups**).

Step 2 Click the Edit icon next to the device or device group for which you want to disable secure store.

Step 3 From the navigation pane, choose **Configure > Security > Secure Store**. The Secure Store Settings window appears, as shown in Figure 9-1.

Step 4 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 or Central Manager from the CLI, use the **cms secure-store clear** EXEC command.

Note When you disable secure store on the primary Central Manager, you should disable secure store on the standby Central manager as well.

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. To modify these properties, choose **Configure > System Properties** from the WAAS Central Manager GUI navigation pane.

Table 9-4 describes the system configuration properties that you can modify.

Table 9-4 Descriptions for System Configuration Properties

System Property	Description
cdm.remoteuser.deletionDaysLimit	Maximum number of days since their last login after which an external user will be deleted from the WAAS Central Manager database. For example, if cdm.remoteuser.deletionDaysLimit is set to 5, an external user 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 1 day. An external user is a user that is defined in an external AAA server and not in the WAAS Central Manager.
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 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.

Table 9-4 Descriptions for System Configuration Properties (continued)

System Property	Description
System.monitoring.collectRate	Rate at which a WAE collects and sends the monitoring report to the WAAS Central Manager (in seconds). 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	WAE 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.
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.print.driverFtpTimeout	Maximum number of seconds to wait for printer driver files to transfer by FTP. The range is 10 to 1800 seconds. The default is 600 seconds.
System.registration.autoActivation	Status of the automatic activation feature, which automatically activates WAE 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.

Table 9-4 Descriptions for System Configuration Properties (continued)

System Property	Description
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.
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 GUI navigation pane, choose **Configure > 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 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 Web Application Filter, page 9-20
- Security Verification, page 9-21

Enabling Web Application Filter

To enable the Web Application Filter, follow these steps:

Step 1 From the WAAS Central Manager GUI, choose **Configure > System Properties**. The Config Properties window appears. (See Figure 9-2.)

Note You cannot enable this feature using the CLI. This feature is disabled by default.

Figure 9-2 Config Properties

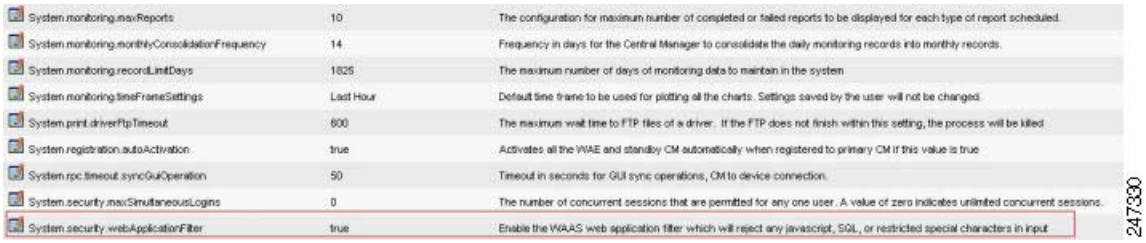

Step 2 Click the Edit icon next to the system.security.webApplicationFilter entry.

The Modifying Config Property window appears.

Step 3 Choose **true** from the Value drop-down list to enable this feature.

A warning appears to advise Central or Device Manager users to log out and then back in after enabling this feature. (See Figure 9-3.)

Figure 9-3 Modifying Config Property

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 9-21
- Sanitization, page 9-21

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. (See Figure 9-4.)

Figure 9-4 Warning

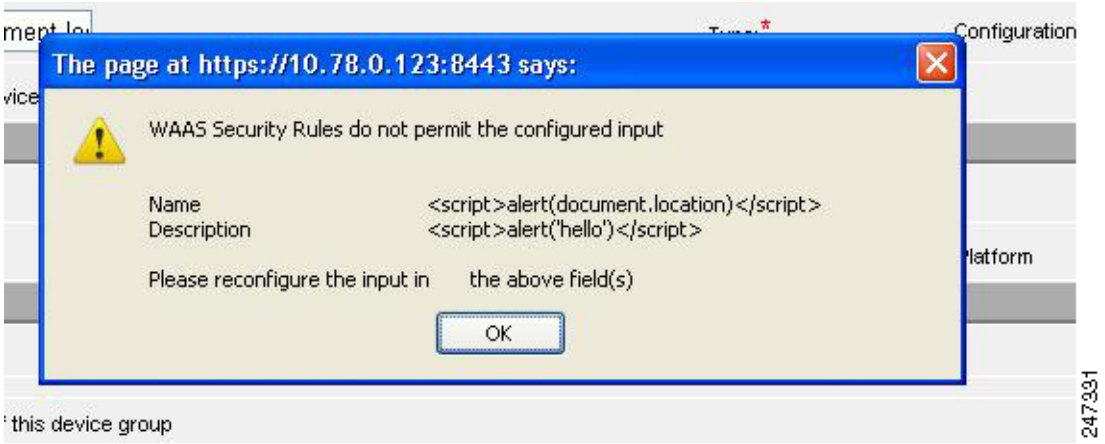

Input data is checked against the following XSS filter rules:

- 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. (See Figure 9-5.)

Figure 9-5 XSS Configuration Data

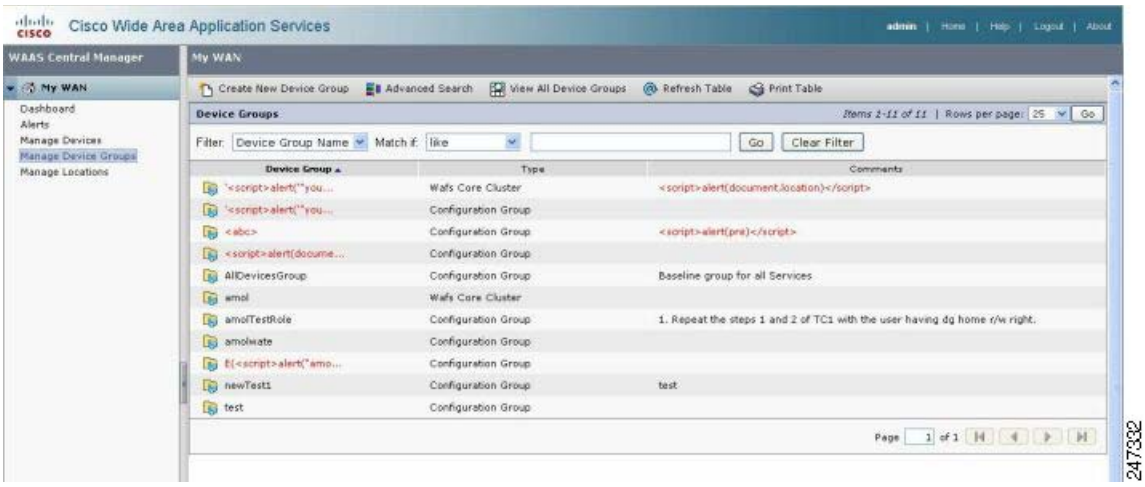

Configuration data coming from the Central Manager that is suspect for XSS is shown in red on the **My WAN > Manage Device Groups > Device Groups** page.

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 9-23 for more information about this feature.)

To configure fast detection of offline WAAS devices, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Configure > 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**.

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**). The Devices (or Device Groups) window appears.

Step 2 Click the **Edit** icon next to the device (or device group) for which you want to configure the alarm overload state.

Step 3 In the navigation pane, choose **Configure > Monitoring > Alarm Overload Detection**. The Alarm Overload Detection Settings window appears.

Step 4 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 5 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 6 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 7 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 16-48.)

To enable email notification, you must configure email server settings for the WAAS Central Manager by following these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**. The Devices window appears.

Step 2 Click the **Edit** icon next to the WAAS Central Manager device for which you want to configure the email server settings.

Note Only SMTP mail servers are supported. If any other type of mail server is configured, the email notification fails.

Step 3 In the navigation pane, choose **Configure > Monitoring > Email Notification Server**. The Configure Email Server Details window appears.

Step 4 In the Mail Server Hostname field, enter the hostname of the SMTP email server that is to be used to send email.

Step 5 In the Mail Server Port field, enter the port number. The default is port 25.

Step 6 In the Server Username field, enter a valid email account username.

Step 7 In the Server Password field, enter the password for the email account.

Step 8 In the From Address field, enter the e-mail address shown as the sender of the email notification.

Step 9 Click **Submit**.

CHAPTER 10

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 3-8.

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, page 10-1
- A Quick Tour of the WAE Device Manager, page 10-2
- WAE Management Workflow, page 10-3
- Managing a Cisco WAE, page 10-4
- Managing a CIFS Accelerator Device, page 10-19
- Managing a WAFS Core Device, page 10-22
- Managing a WAFS Edge Device, page 10-23
- Monitoring the WAE, page 10-23
- Viewing WAE Logs, page 10-32

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.

To launch the WAE Device Manager, use one of the following methods:

- 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.

- From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**, click the **Edit** icon next to the device that you want to manage, and click the **Device GUI** button at the bottom of the window.

The WAE Device Manager interface appears. (See Figure 10-1.)

Figure 10-1 WAE Device Manager Interface

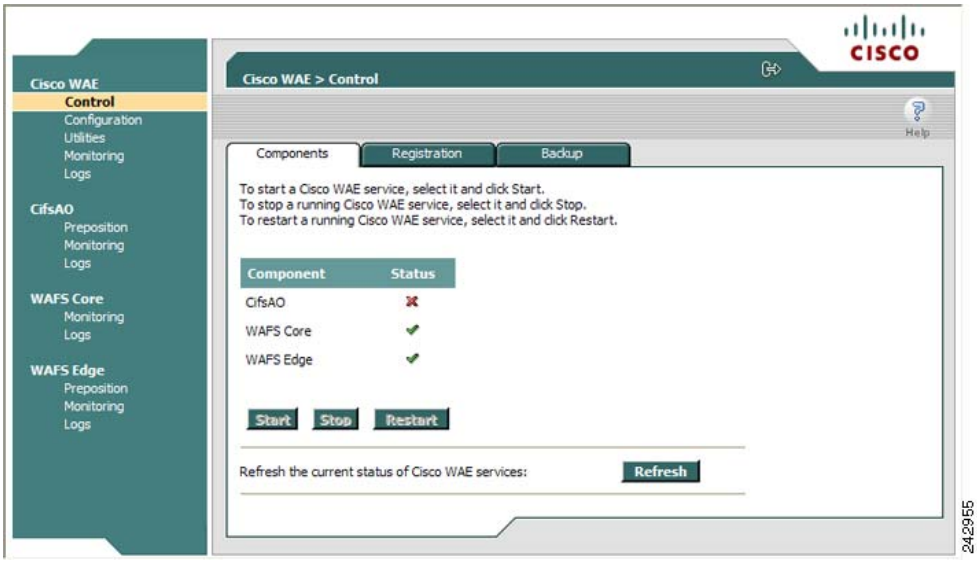

A Quick Tour of the WAE Device Manager

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 10-4.
- CifsAO—Allows you to monitor preposition tasks, view WAFS device statistics, and view the log. For more information, see the "Managing a CIFS Accelerator Device" section on page 10-19.

  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 12-2.
- WAFS Core—Allows you to monitor WAFS Core statistics and view log information. For more information, see the "Managing a WAFS Core Device" section on page 10-22.

The WAFS Core option only appears if you have configured this WAAS device as a core device in legacy WAFS mode. For more information, see Chapter 11, "Configuring Wide Area File Services."

- WAFS Edge—Allows you to monitor preposition tasks, monitor WAFS Edge statistics, and view log information. For more information, see the "Managing a WAFS Edge Device" section on page 10-23.

  The WAFS Edge option appears only if you have configured this WAAS device as an edge device in legacy WAFS mode. For more information, see Chapter 11, "Configuring Wide Area File Services."

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 10-5.
- Register and unregister the WAE as described in the "Registering and Unregistering a WAE" section on page 10-6.
- Back up and restore configuration files as described in the "Backing Up the Configuration Files" section on page 10-7.
- Configure Windows authentication as described in the "Configuring Windows Authentication" section on page 10-10.
- Define component-specific notification recipients as described in the "Defining Notification Settings" section on page 10-15.
- Run WAE maintenance utilities as described in the "Utilities Option" section on page 10-17.
- View the details, current status, and history of preposition tasks performed on WAFS devices as described in the "Preposition Option" section on page 10-20.
- View SNMP-generated information and graphs about each WAE component as described in the "Monitoring the WAE" section on page 10-23.
- View the logs for each WAE component as described in the "Viewing WAE Logs" section on page 10-32.

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 10-2 shows the Cisco WAE Control window.

Figure 10-2 Cisco WAE Control Window

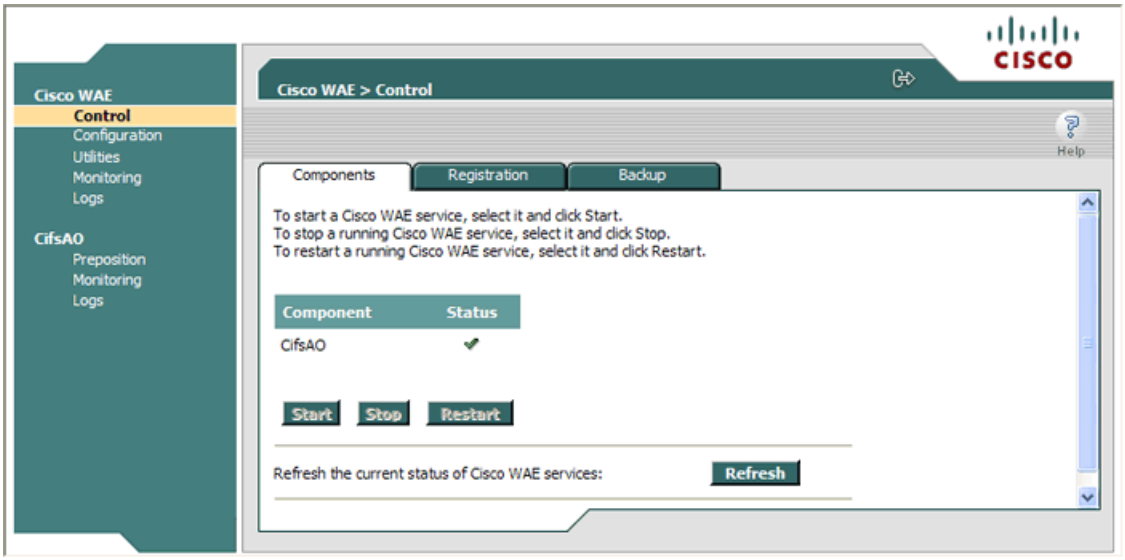

The Cisco WAE menu item includes the following options:

- Control—Enables you to control the WAE and its components as described in the "Control Option" section on page 10-4.
- Configuration—Enables you to perform basic configuration tasks as described in the "Configuration Option" section on page 10-8.
- Utilities—Enables you to run various maintenance utilities on the WAE as described in the "Utilities Option" section on page 10-17.
- 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 10-23.
- Logs—Enables you to view event logs for various WAE subsystems as described in the "Viewing WAE Logs" section on page 10-32.

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 10-5.
- 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 10-6.

- 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 10-7 and the "Restoring the Configuration Files" section on page 10-7.

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 10-3.

Figure 10-3 Components Tab—Starting Components

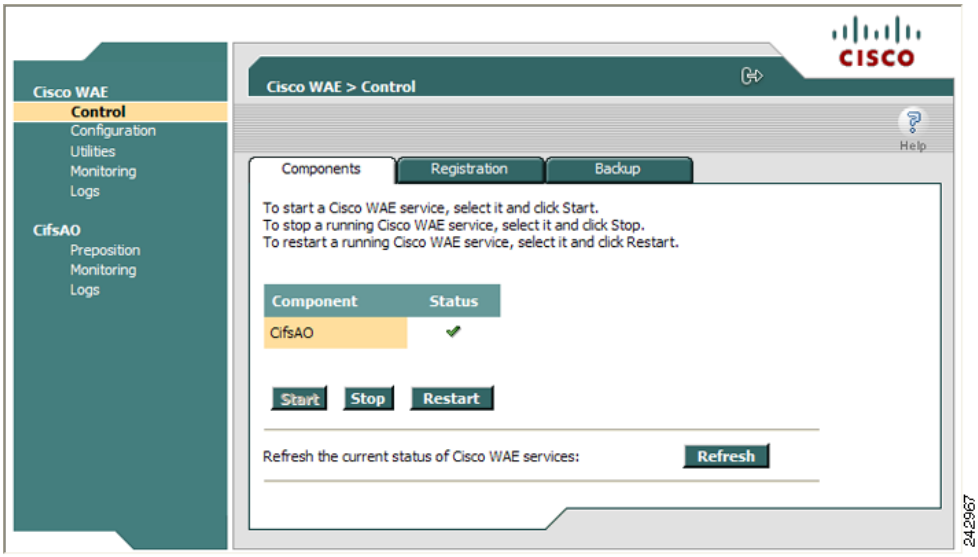

- To stop a component, choose the component from the list and click **Stop**.

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 10-4.)

Figure 10-4 Cisco WAE Control—Registration Tab

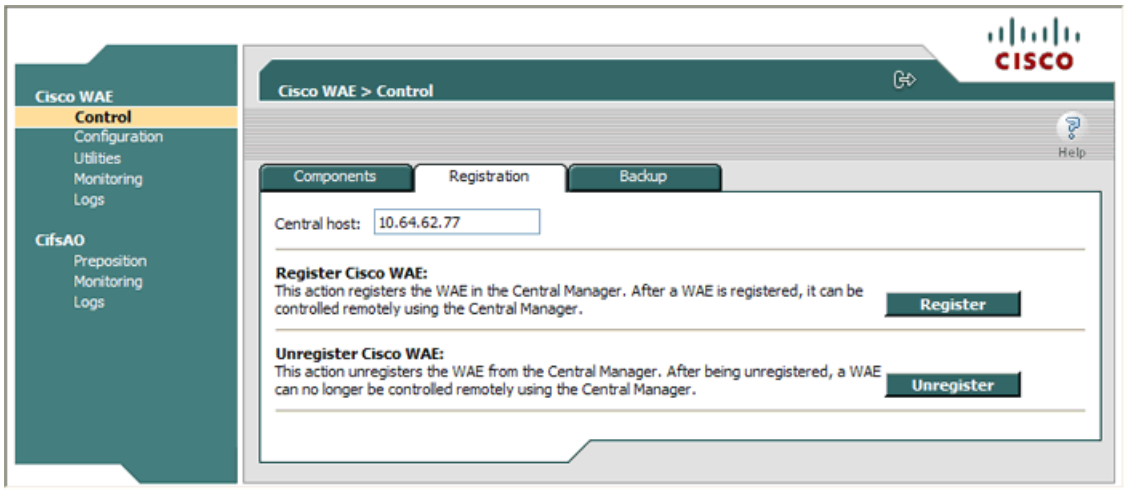

Step 2 In the Central Host field, verify that the address of the WAAS Central Manager is displayed. If no address appears in this field, then the WAE is not registered with a Central Manager.

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 10-5.)

Figure 10-5 Cisco WAE Control —Backup Tab

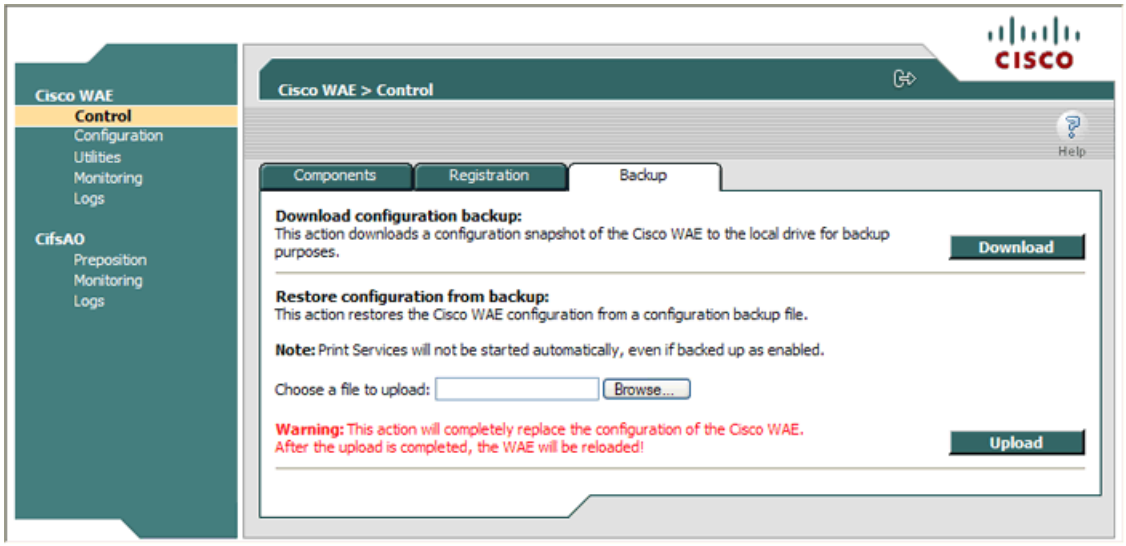

Step 2 In the Download configuration backup area, click **Download**.

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 10-7.

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 10-8.
- 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 10-9.
- Windows Authentication—Allows you to define the settings required by the WAE for Windows authentication to enable device login, disconnected mode, and CLI configuration. For more information, see the "Configuring Windows Authentication" section on page 10-10.
- Print Services—Allows you to configure print services. For more information on configuring print services, see Chapter 13, "Configuring and Managing WAAS Legacy Print Services."
- 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 10-15.

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 10-6.)

Figure 10-6 WAE Configuration—SNMP Tab

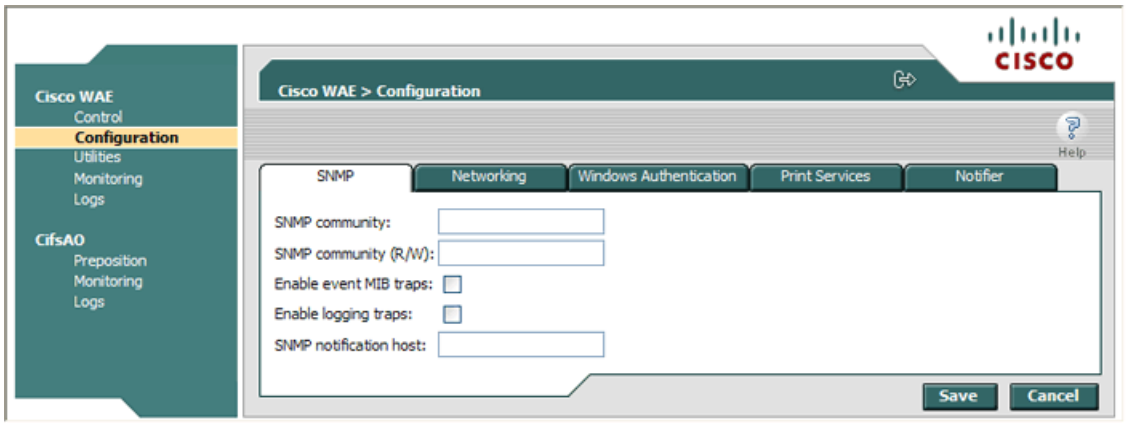

This tab allows you to configure the following settings:

- 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 10-7) 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 10-7 Cisco WAE Configuration—Networking Tab

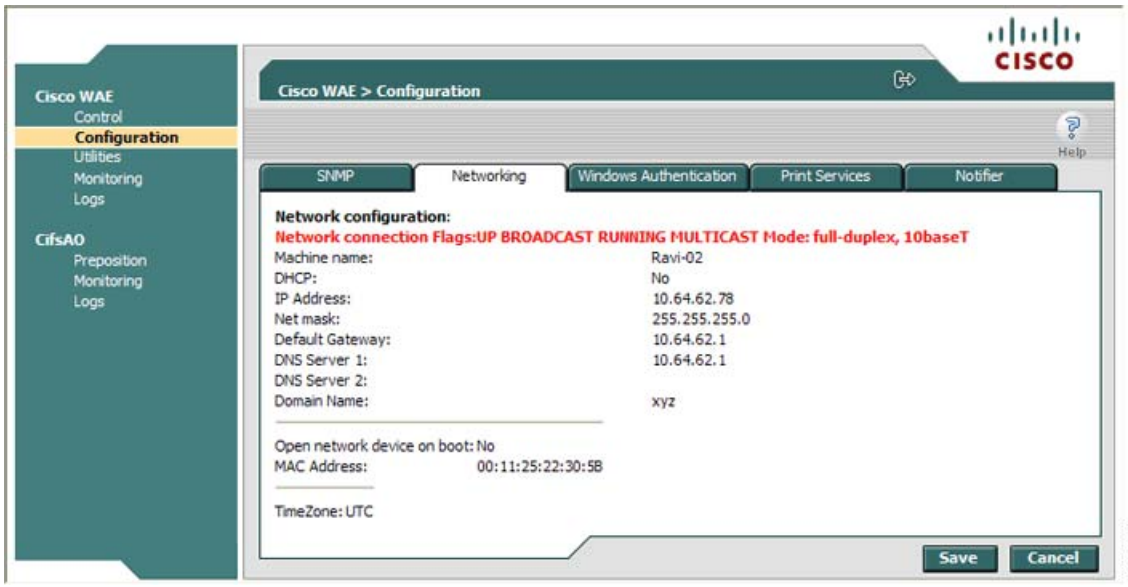

The Networking tab contains the following information:

- 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, page 10-10
- Supported Authentication Methods, page 10-10
- LDAP Server Signing, page 10-11
- Setting Up Windows Authentication, page 10-11
- Checking the Status of Windows Authentication, page 10-13

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 (WAFS 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.

WAFS 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 **s**ettings. 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 both print services' and 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 "start_tls"** 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 10-8.)

Figure 10-8 Cisco WAE Configuration—Windows Authentication Tab

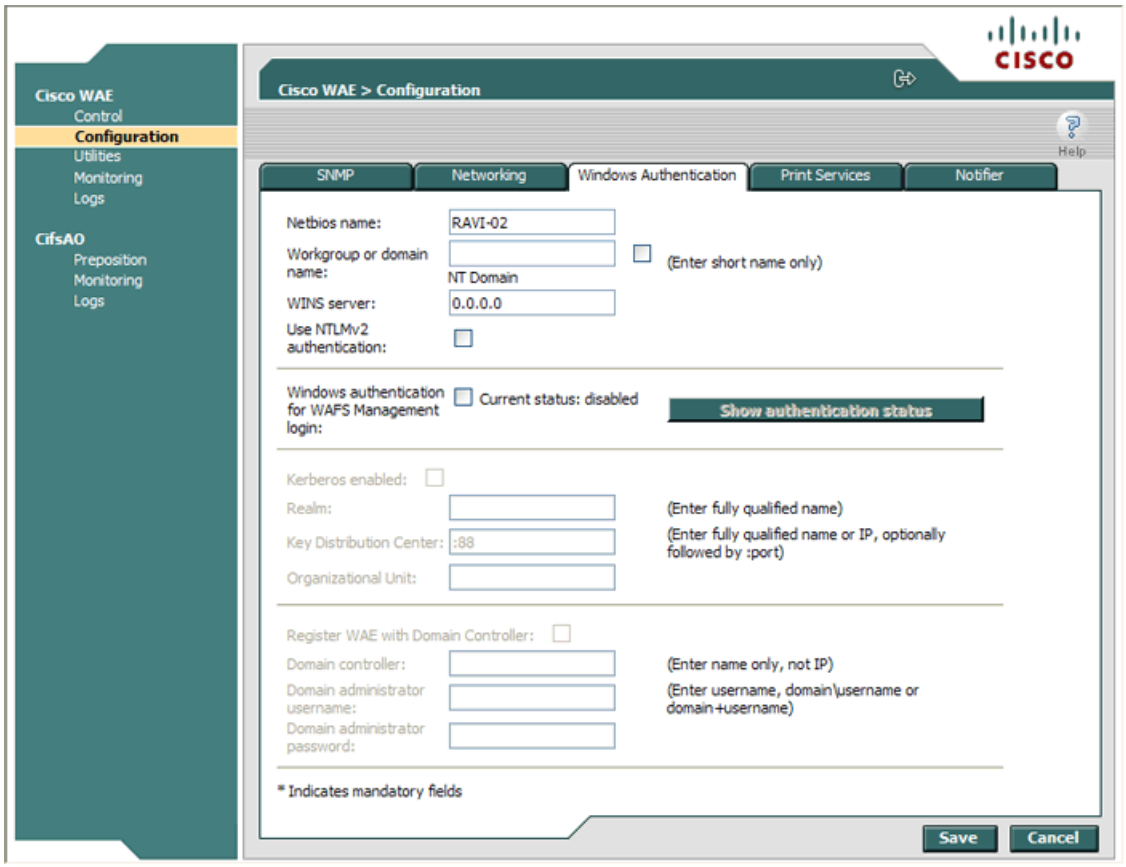

Step 3 Enter the NetBIOS name.

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 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 "start_tls"** 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 10-13.

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 10-11. 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 10-9.)

Figure 10-9 Authentication Details Window

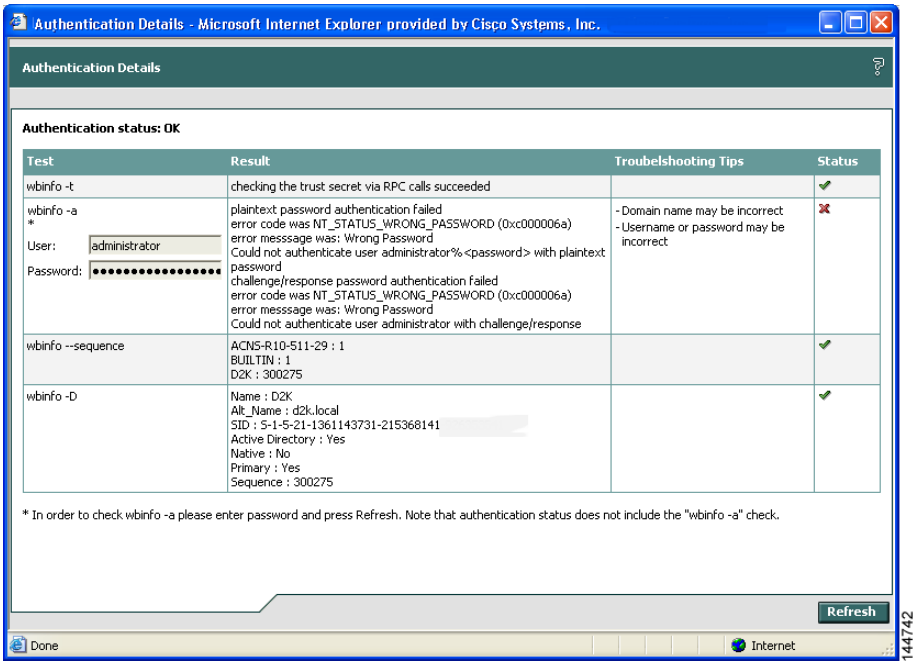

Step 3 Check the Authentication status field at the top of the window.

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 10-1 describes these tests.

Table 10-1 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 10-10.)

Figure 10-10 Notifier Tab

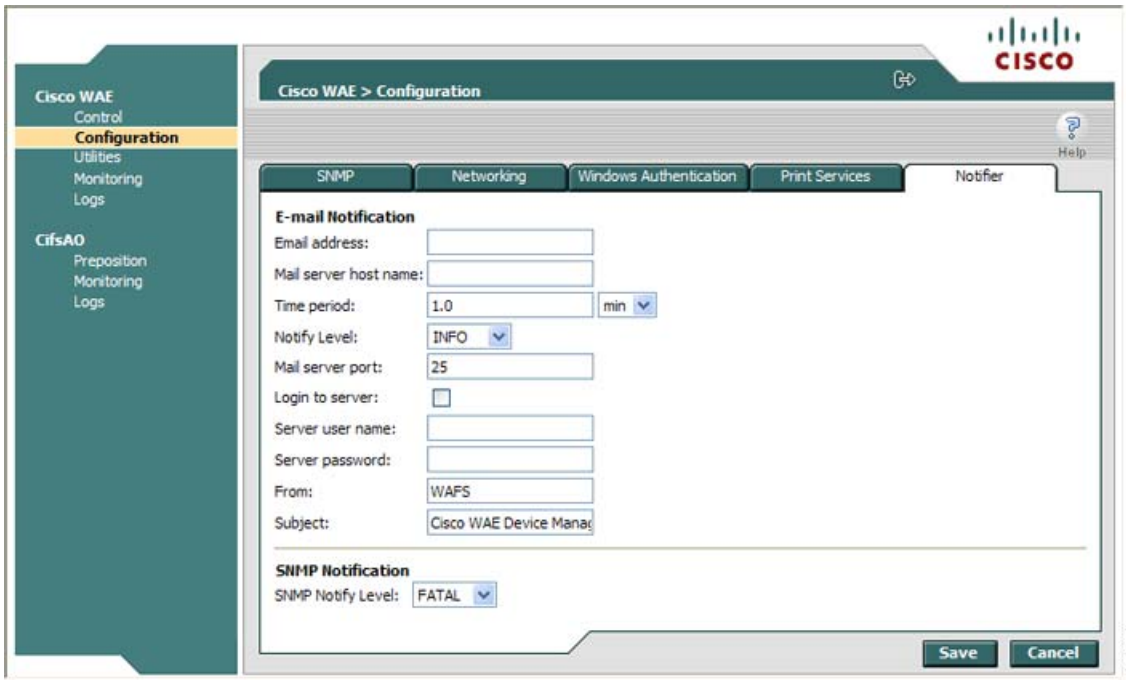

Step 2 In the Email address field, enter the address to which notifications about this WAE are sent.

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 10-17.
- WAFS Cache Cleanup—Allows you to remove all files from the WAFS cache. For more information, see the "Running the Cache Cleanup Utility" section on page 10-18.
- File Server Rename—Allows you to rename a file server in the WAFS cache. For more information, see the "Running the File Server Rename Utility" section on page 10-19.

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 10-11.)

Figure 10-11 Utilities—Support Tab

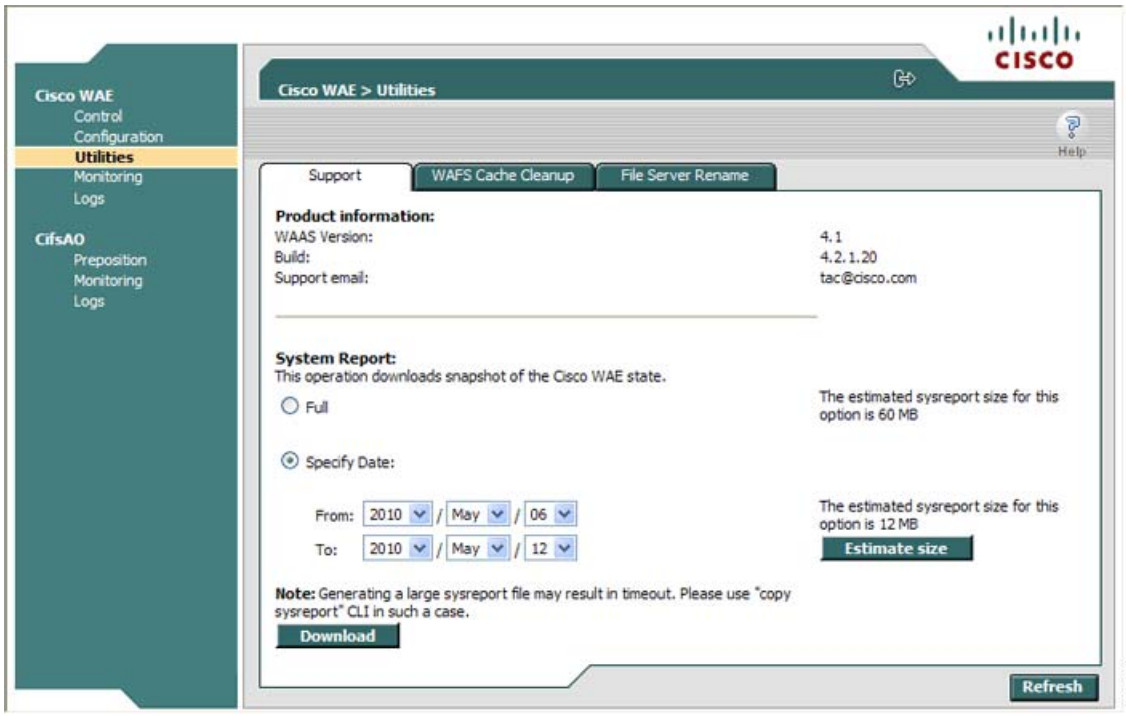

Step 2 In the System Report area, choose one of the following radio buttons:

- **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 WAFS 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 10-12.)

Figure 10-12 Utilities—WAFS Cache Cleanup Tab

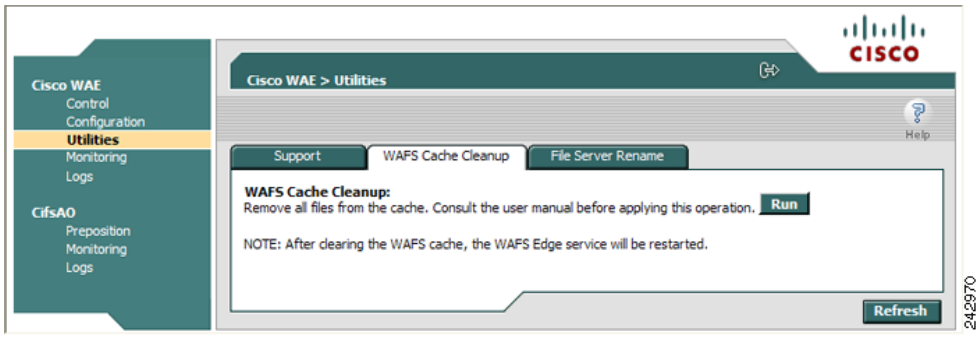

Step 2 Click **Run** to erase the contents of the cache.

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 WAFS cache.

To run the file server rename utility, follow these steps:

Step 1 If the WAFS Edge component is running, stop it as described in the "Starting and Stopping Components" section on page 10-5.

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 WAFS 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 10-20.

- Monitoring—Allows you to view WAFS device statistics in tables and graphs as described in the "Monitoring the Cisco WAE Component" section on page 10-25.
- 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 10-34.

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 10-13.)

Figure 10-13 CifsAO > Preposition Window

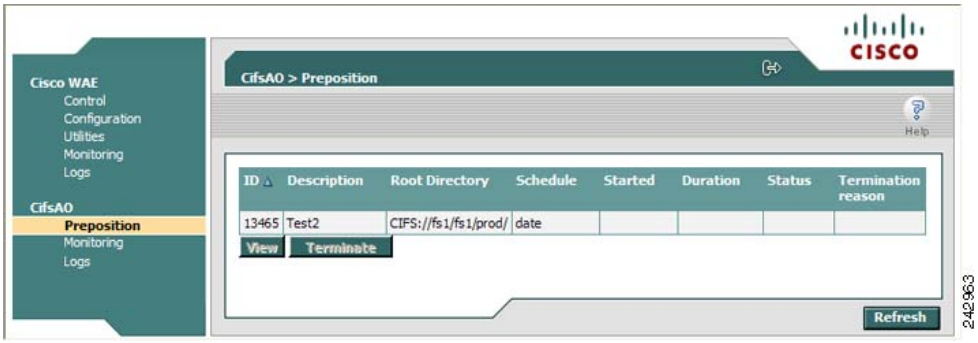

The Preposition window contains a table that displays all the preposition policies assigned to this WAFS Edge device. For each policy, the following information is displayed:

- 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 10-14.)

Figure 10-14 Preposition Task Details

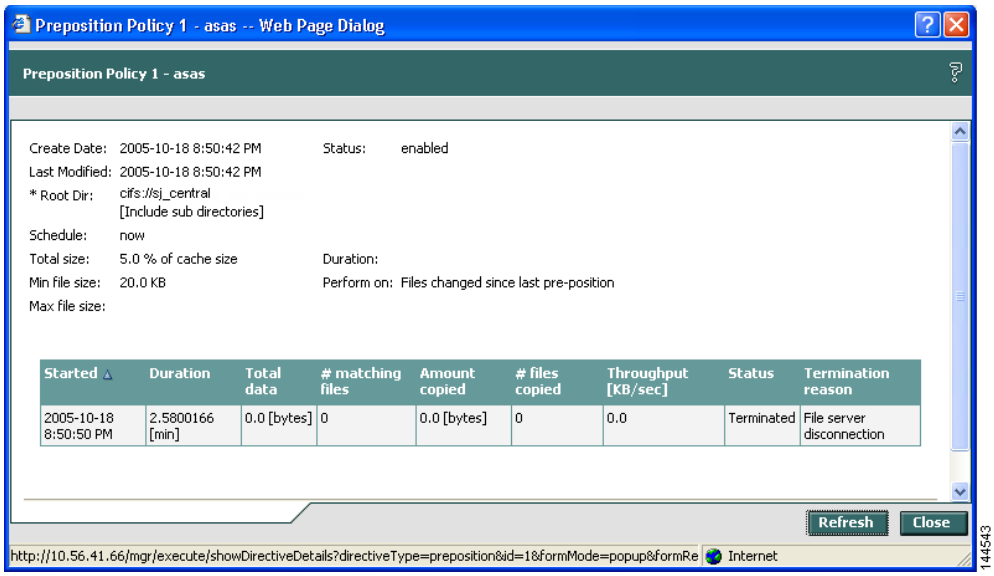

The top half of the Preposition Policy window displays the following details about the selected policy:

- 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.

Managing a WAFS Core Device

The WAFS Core option in the navigation area allows you to monitor WAFS Core statistics and view log information. The WAFS Core option appears only if you are using WAFS legacy mode.

Note Legacy mode WAFS is deprecated in WAAS version 4.2.1. It still operates but is not supported and will be removed in a future version. If you enable legacy mode WAFS, an alarm is raised on the Central Manager and you are alerted in all Central Manager GUI pages and the CLI if you configure any legacy WAFS settings. Legacy WAFS users should migrate to the transparent CIFS accelerator.

The WAFS Core option includes the following menu items:

- Monitoring—Allows you to view WAFS Core statistics in tables and graphs as described in the "Monitoring the Cisco WAE Component" section on page 10-25.
- Logs—Allows you to view the event log related to the WAFS Core. For more information, see the "Viewing Cisco WAE Logs" section on page 10-34.

Managing a WAFS Edge Device

The WAFS Edge menu item in the navigation area allows you to monitor preposition tasks, monitor WAFS Edge statistics, and view log information. The WAFS Edge option appears only if you are using WAFS legacy mode.

Note Legacy mode WAFS is deprecated in WAAS version 4.2.1. It still operates but is not supported and will be removed in a future version. If you enable legacy mode WAFS, an alarm is raised on the Central Manager and you are alerted in all Central Manager GUI pages and the CLI if you configure any legacy WAFS settings. Legacy WAFS users should migrate to the transparent CIFS accelerator.

The WAFS Edge 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 10-20. Preposition works the same for WAFS legacy mode and transparent CIFS accelerator mode.
- Monitoring—Allows you to view WAFS Edge device statistics in tables and graphs as described in the "Monitoring the Cisco WAE Component" section on page 10-25.
- Logs—Allows you to view the event log related to the WAFS Edge device. For more information, see the "Viewing Cisco WAE Logs" section on page 10-34.

Monitoring the WAE

The Monitoring option available for the Cisco WAE, WAFS Core, WAFS Edge, 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 10-2.

Table 10-2 Monitoring Options by Component

Component	Monitored Statistics
Cisco WAE	CPU and disk drive utilization
Transparent CIFS accelerator	CIFS traffic and cache
WAFS Core	Connectivity
WAFS Edge	Connectivity, CIFS traffic, and cache

This section contains the following topics:

- Monitoring Graphs, page 10-24
- Monitoring the Cisco WAE Component, page 10-25

- Monitoring a WAFS Core, page 10-26
- Monitoring a Transparent CIFS Accelerator or a WAFS Edge Device, page 10-28

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 10-15 shows a sample screen when a user chooses to view the index graphs.

Figure 10-15 Sample Index Graph Window

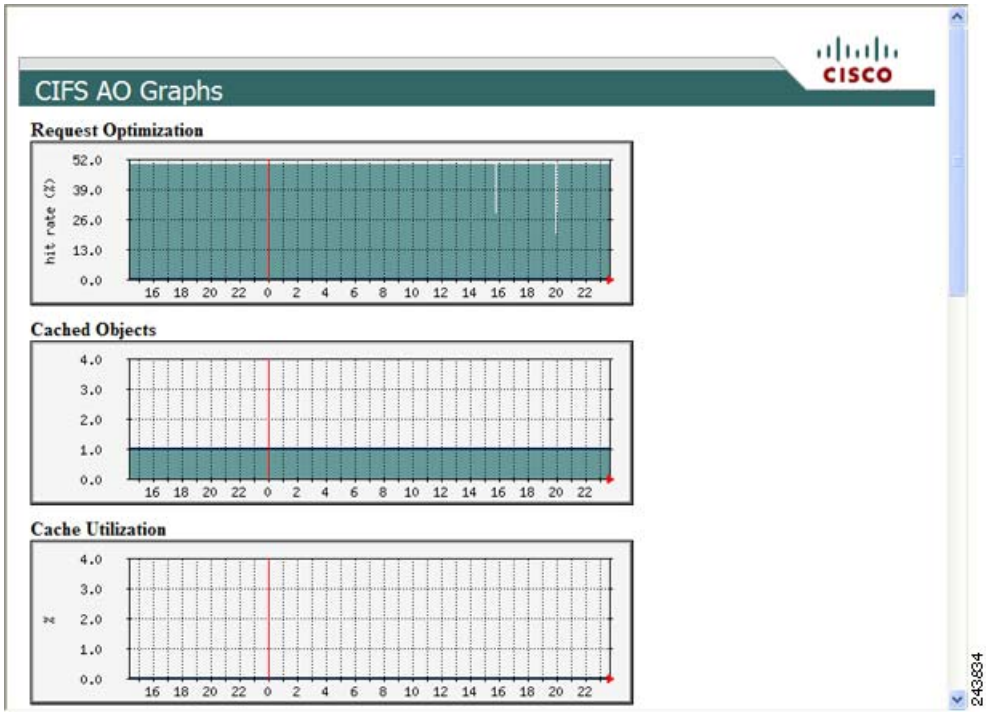

Tip Each graph in an index window acts as a link. Clicking on the graph displays all four historical graphs for the selected statistic. For example, clicking the Request Optimization graph in the index graphs window displays the daily, weekly, monthly and yearly Request Optimization historical graphs. Clicking the Back button in the browser returns you to the index graphs.

Figure 10-16 shows a sample screen when a user chooses to view the historical graphs for a particular statistic.

Figure 10-16 Sample Historical Graph Window

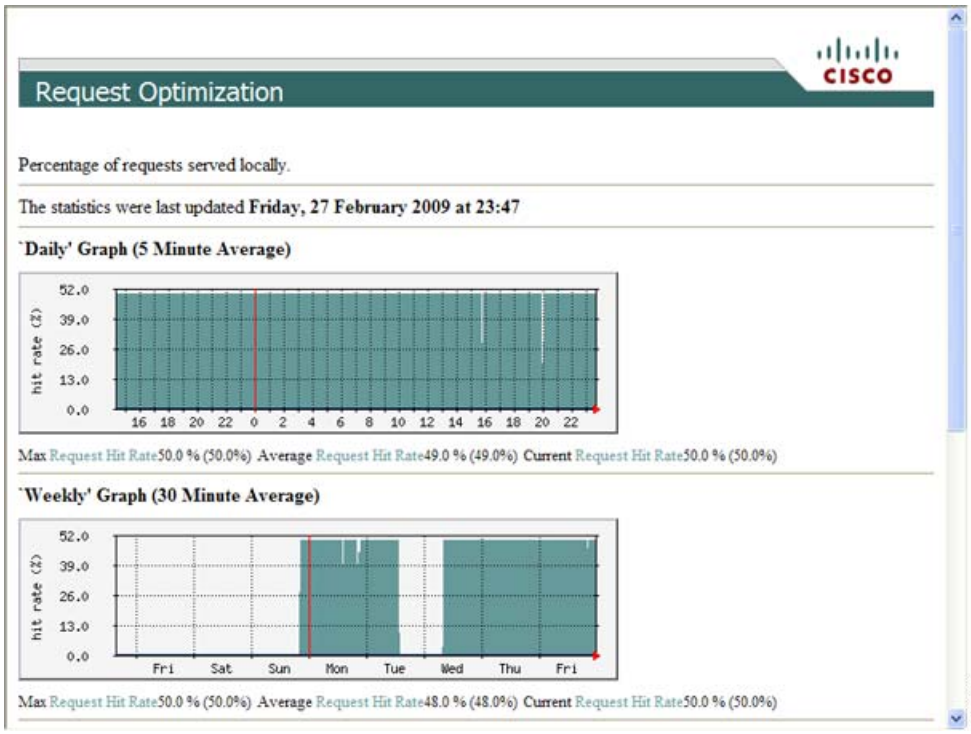

Note Graphs can be printed using the Print command in your browser.

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 10-17.)

Figure 10-17 Cisco WAE Monitoring Window

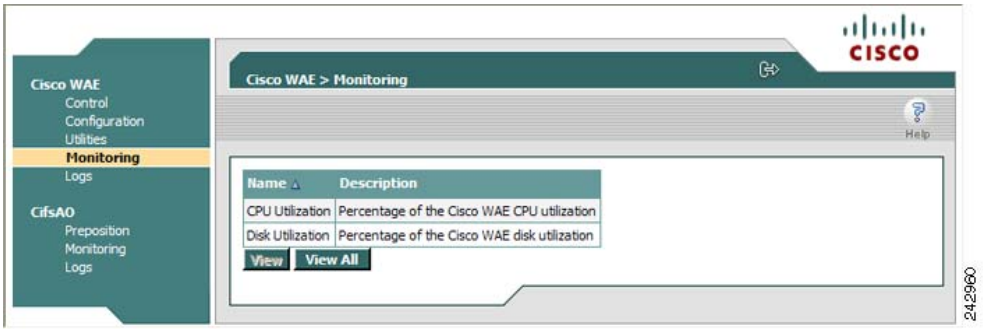

Step 2 Do one of the following:

- 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 WAFS Core

The Monitoring option for the WAFS Core menu item displays the following two tabs:

- Connectivity—Displays a table of connectivity statistics for the WAFS Core. This table is displayed by default.
- Graphs—Displays a list of graphs that are available for the WAFS Core.

To monitor a WAFS Core, follow these steps:

Step 1 In the navigation area, click **Monitoring** under the WAFS **Core** component.

The Connectivity tab on the WAFS Core Monitoring window appears. (See Figure 10-18.)

Figure 10-18 WAFS Core Monitoring—Connectivity Tab

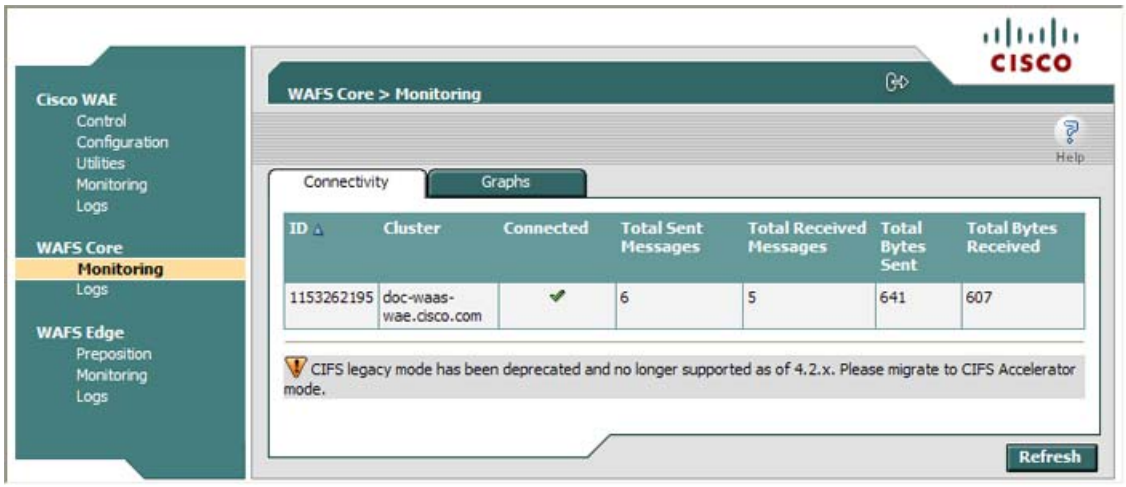

The Connectivity tab displays a table that contains the following data about the WAFS Core:

- ID—Alphanumeric system identifier for the WAFS Core.
- Cluster—Name of the core cluster to which this WAFS Core belongs, if any.
- Connected—Indicator of whether the WAFS Core is currently connected to (✔) or disconnected from (✖) its WAFS Edge devices.
- Total Sent Messages—Total number of messages sent from this WAFS Core since activation.
- Total Received Messages—Total number of messages received by this WAFS Core since activation.
- Total Bytes Sent—Total number of bytes sent from this WAFS Core since activation.
- Total Bytes Received—Total number of bytes received from this WAFS Core since activation.

Step 2 Click the **Graphs** tab. (See Figure 10-19.)

Figure 10-19 WAFS Core Monitoring – Graphs Tab

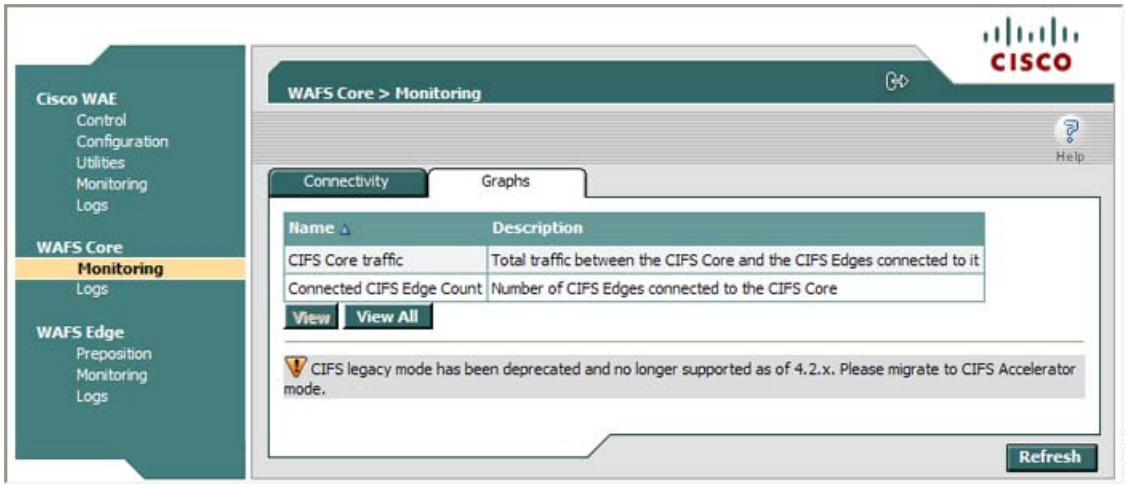

The following historical graphs are available for a WAFS Core component:

- Connected WAFS Edge counts—Number of WAFS Edge devices currently connected to the selected WAFS Core. This graph is useful for detecting WAFS Edge device disconnections.
- WAFS Core traffic—Total volume of traffic (in kilobits) between the WAFS Core and each of the WAFS Edge devices connected to it. The green line represents transmitted traffic; the blue line represents received traffic.

Step 3 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 WAFS Core component.

Monitoring a Transparent CIFS Accelerator or a WAFS Edge Device

Monitoring a transparent CIFS accelerator device or a WAFS Edge device is almost the same, except that the transparent CIFS accelerator does not monitor connectivity.

The Monitoring option displays the following tabs:

- Connectivity—Displays a table of connectivity statistics for the WAFS Edge device. This tab does not appear when using transparent CIFS accelerator mode.
- 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 or WAFS Edge device, follow these steps:

Step 1 In the navigation area, click **Monitoring** under the **CifsAO** or **WAFS Edge** menu.

The Monitoring window appears. (See Figure 10-20.) If you are using transparent CIFS accelerator mode, the Connectivity tab is not used and the CIFS tab is displayed first; skip to Step 2.

Figure 10-20 WAFS Edge Monitoring—Connectivity Tab

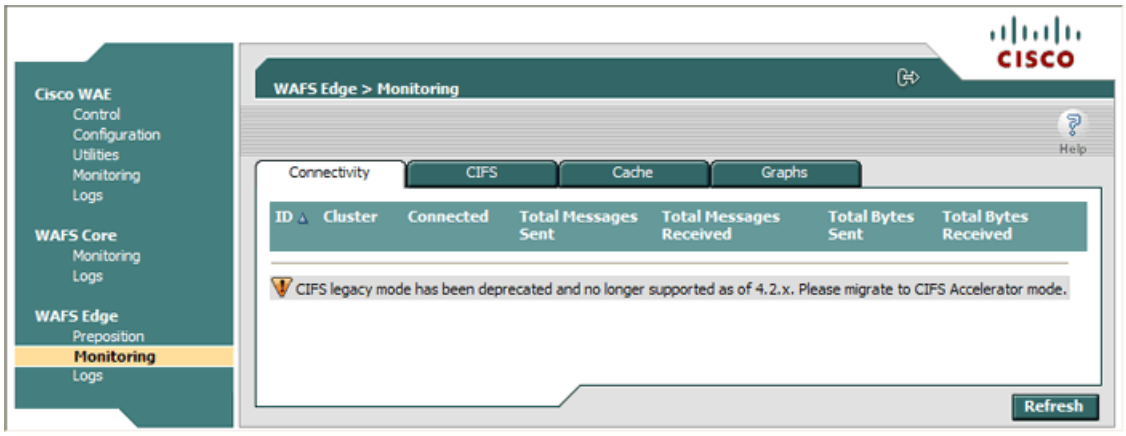

The Connectivity tab displays a table that contains the following data about the WAFS Edge device:

- ID—Alphanumeric system identifier for the WAFS Edge device.
- Cluster—Name of the Core cluster to which this WAFS Edge device is connected, if any.
- Connected—Indicator of whether the WAFS Edge device is currently connected to (✔) or disconnected from (✖) the WAFS Core.
- Total Messages Sent —Total number of messages sent from this WAFS Edge device since activation.
- Total Messages Received—Total number of messages received by this WAFS Edge device since activation.
- Total Bytes Sent—Total number of bytes sent from this WAFS Edge device since activation.
- Total Bytes Received—Total number of bytes received by this WAFS Edge device since activation.

Step 2 Click the **CIFS** tab.

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 3 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 WAFS device cache.
- Current cache disk usage—Current amount of disk space (in kilobytes) used by the WAFS 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 WAFS device cache.
- Current cache resources—Current number of resources contained in the WAFS 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 WAFS device to begin evicting resources.
- Cache size low watermark—Percentage of disk usage that causes the WAFS device to stop evicting resources.
- Cache resources high watermark—Percentage of total cache resources that causes the WAFS device to begin evicting resources.
- Cache resources low watermark—Percentage of total cache resources that causes the WAFS device to stop evicting resources.
- Last evicted resource age—Amount of time that the last-evicted resource spent in the WAFS device cache.
- Last evicted resource access time—Last time that the last-evicted resource was accessed.

Step 4 Click the **Graphs** tab (see Figure 10-21).

Figure 10-21 WAFS Edge Monitoring – Graphs Tab

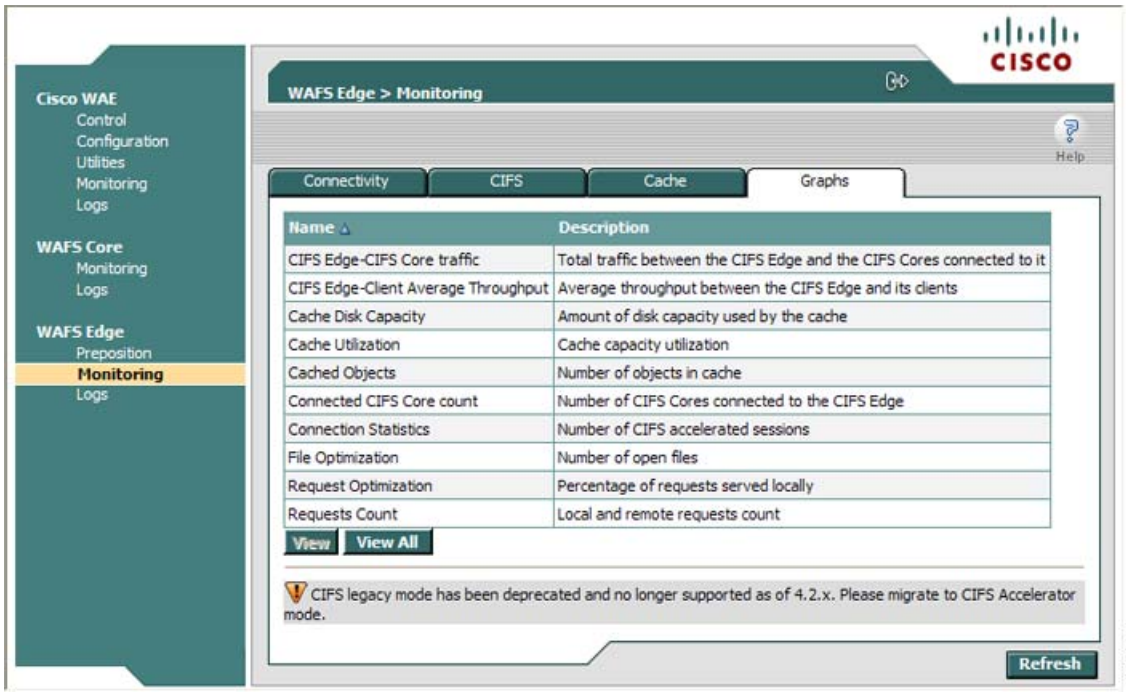

The following historical graphs are available for the device:

- CIFS Edge - CIFS Core traffic—Total volume of traffic (in kilobytes) between the WAFS Edge device and each of the WAFS Cores connected to it. The green line represents transmitted traffic; the blue line represents received traffic. This item appears only if you are using WAFS legacy mode.
- CIFS Edge - Client Average Throughput—Total volume of traffic between the WAFS Edge device and the clients it serves, divided by total uptime (including idle time). Displayed in kilobytes per second. This item appears only if you are using WAFS legacy mode.
- Cache Disk Capacity—Amount of disk space (in megabytes) used by the WAFS device cache.
- Cache Utilization—Percentage of disk space and percentage of resources used by the cache, based on defined limits.
- Cached Objects—Total number of objects (files and directories) contained in the cache.
- Client Average Throughput—Total volume of traffic between the WAFS Edge device and the clients it serves, divided by total uptime (including idle time). Displayed in kilobytes per second. This item appears only if you are using CIFS transparent accelerator mode.
- Connected CIFS Core count—Number of WAFS Cores connected to the selected WAFS Edge device. This item appears only if you are using WAFS legacy mode.

Note The WAFS Edge device can be connected to multiple WAFS Cores for potential higher availability.

- Connection Statistics—Number of accelerated CIFS sessions on the device.
- File Optimization—Total number of open CIFS files.

- Request Optimization—Percentage of user requests answered by the cache (as opposed to sending a request remotely over the WAN to the file server for the user-requested content).
- Requests Count—Average rate of requests handled locally (client requests answered through the cache) and average rate of requests handled remotely (client requests answered through the remote file server). The request count is displayed in requests per second.

Step 5 Do one of the following actions:

- Choose a graph in the table and then click **View** to display a popup window with all four historical graphs for the selected statistic.
- Click **View All** to display the index window with the daily graphs for the WAFS Edge device.

Viewing WAE Logs

You can view event information logged by the Cisco WAE, CifsAO, WAFS Core, and WAFS Edge components. The event information available varies based on the component that you are viewing.

This section contains the following topics:

- WAE Logs, page 10-32
- Viewing Cisco WAE Logs, page 10-34

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 10-32
- Viewing Log Entries, page 10-33
- Saving Log File Information, page 10-33

Setting Display Criteria

All WAE logs allow you to set the criteria for the data that you want to display as shown in Figure 10-22.

Figure 10-22 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 10-32.)

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 10-32.

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 WAFS Core generates one log that displays all events related to WAFS Core operation. The WAFS Edge device also generates one log that displays all events related to the WAFS Edge operation. If you are using transparent CIFS accelerator mode, the CIFS accelerator generates one log that displays all events related to CIFS accelerator operation.

To view Cisco WAE, CIFS accelerator, WAFS Core, or WAFS Edge logs, follow these steps:

Step 1 In the navigation area, click the **Logs** option under the Cisco WAE, WAFS Core, WAFS Edge, or CifsAO component.

Figure 10-23 shows the Logs window for the Cisco WAE component.

Figure 10-23 Cisco WAE Component Logs Window

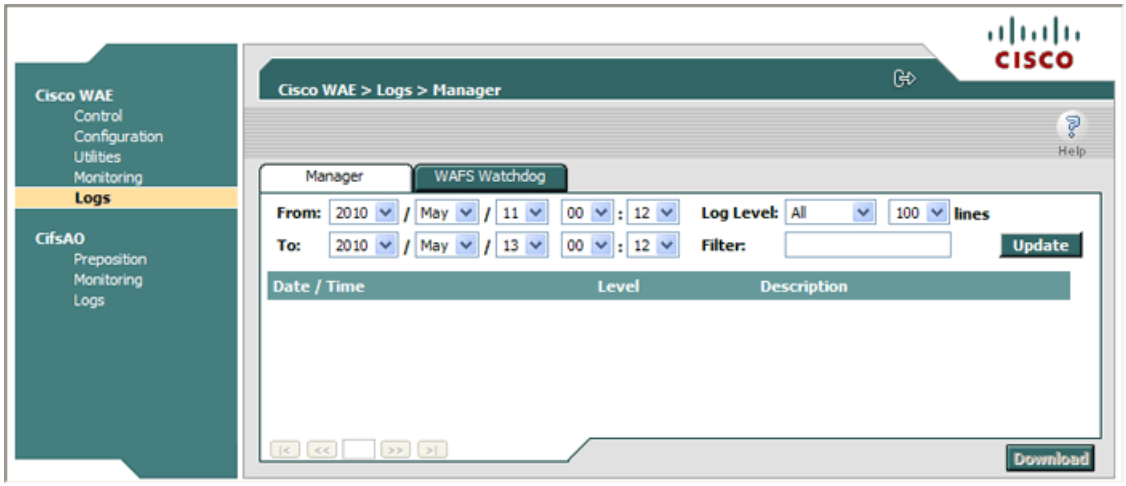

Step 2 If you selected the Cisco WAE, click the **Manager** or **WAFS Watchdog** tab to choose the log that you want to view.

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 10-32.)

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.

PART 3

Configuring WAAS Services

CHAPTER 11

Configuring Wide Area File Services

This chapter describes how to configure Wide Area File Services (WAFS), which allows branch office users to more efficiently access data stored at centralized data centers. The WAFS feature overcomes the WAN latency and bandwidth limitations by caching data on Edge WAEs near branch office users.

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, page 11-1
- Overview of File Services Features, page 11-3
- Preparing for File Services, page 11-8
- Configuring File Services, page 11-9
- Managing File Services, page 11-33

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. 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 11-1 shows a typical deployment scenario after WAAS file services have been set up.

Figure 11-1 WAAS File Services Solution

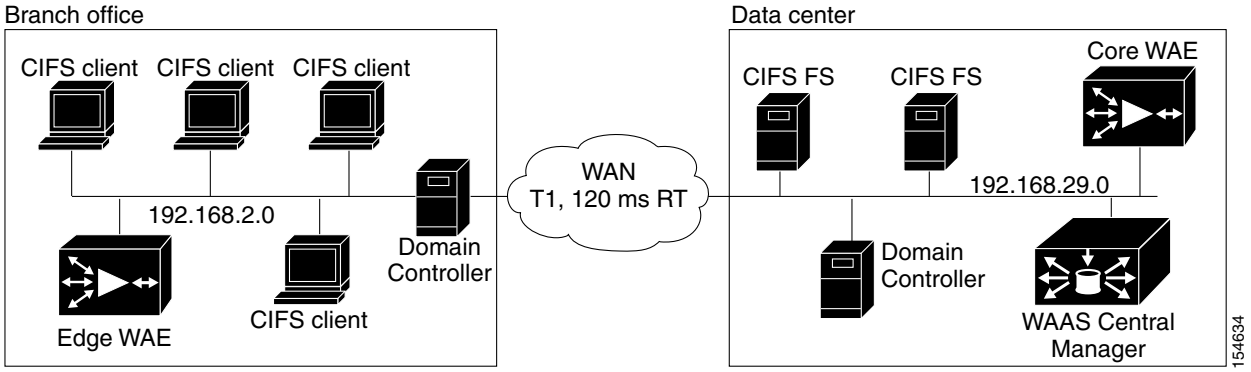

Overview of File Services Features

This section provides an overview of the WAAS file services features and contains the following topics:

- Automatic Discovery, page 11-4
- Prepositioning, page 11-4
- Data Coherency, page 11-5
- Data Concurrency, page 11-6
- Microsoft Interoperability, page 11-7

WAAS version 4.1.1 and later includes two mutually exclusive WAFS modes:

- Transparent CIFS accelerator mode—This mode was introduced in WAAS version 4.1.1. The new transparent mode requires no core, edge, or connectivity configuration. It relies on automatic discovery and transparently accelerates CIFS traffic with no configuration needed. Disconnected mode is not supported. We recommend this mode for ease of use.
- Legacy mode—This mode, present in all WAAS versions, requires that you configure a Core cluster, Edge file services, and connections between the Core and Edge devices. Use this mode if you need to interoperate with WAAS 4.0.x devices or if disconnected mode operation is essential.

  This mode also includes the ability to automatically discover file servers, but there are some limitations. For more information, see the "Automatic Discovery with Legacy Mode" section on page 11-4.

Note Legacy mode WAFS is deprecated in WAAS version 4.2.1. It still operates but is not supported and will be removed in a future version. If you enable legacy mode WAFS, an alarm is raised on the device and you are alerted in all Central Manager GUI pages and the CLI if you configure any legacy WAFS settings. Legacy WAFS users should migrate to the transparent CIFS accelerator.

If you enable the transparent CIFS accelerator on a WAE device or group, the system disables any existing Core or Edge configurations and the Core and Edge configuration windows; additionally, the device cannot belong to a Core cluster. In the CIFS accelerator mode, the WAE device acts as an Edge and a Core simultaneously, depending on the direction of each request. There is no need for a manual setting.

Note WAFS acceleration operates only if the peer WAEs at each end of the connection use the same WAFS mode (transparent CIFS accelerator or legacy).

Automatic Discovery

The automatic discovery feature allows you to enable WAFS without having to register individual file servers in the WAAS Central Manager. With the automatic discovery feature, WAAS will attempt to automatically discover and connect to a new file server when a transparent mode CIFS request is received.

Automatic Discovery with Legacy Mode

The legacy mode of operation also includes the ability to automatically discover file servers, with some limitations. The automatic discovery feature operates by default for CIFS requests to unregistered file servers when you are using legacy mode.

If there are multiple paths to the file server, WAAS chooses the path with the lowest latency. If the latency between the core WAE and the discovered server is more than 25 milliseconds, the server is considered to be too far away, and the connection will not be optimized. Additionally, if the latency between the edge WAE and the server is less than 2 milliseconds, the server is considered to be local, and the connection will not be optimized.

Note If you are using the WAFS legacy mode, and relying on automatic discovery instead of registering file servers, the following WAFS features are not available: prepositioning, dynamic shares, and disconnected mode. Additionally, if the file server requires a digital signature, WAFS cannot cache its data.

As new file servers are added to the policy engine dynamic map by this mode of automatic discovery, you can see them by using the **show policy-engine application dynamic** EXEC command. A file server remains in the dynamic map for three minutes after the last connection to it is closed.

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.

Note Prepositioning includes the ability to configure multiple roots. See the "Creating a New Preposition Directive" section on page 11-27.

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.

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 11-7)

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 Edge WAE 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 edge 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 edge devices on which you want to enable file services. An edge device may also serve as a core device if it also exports local file servers to other edge devices.
- Identify the file servers that you want to export, and refer to Table 11-1 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 WAAS file services feature does not support file servers that use the FAT32 file system. You can use the policy engine rules to exclude any file servers that use the FAT32 file system from CIFS optimization.

Table 11-1 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 11-1 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 Vista clients, WAAS 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).

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 the following amounts of memory in order to support file services:

- Transparent CIFS accelerator mode—1 GB RAM
- Legacy mode, edge file services—1 GB RAM
- Legacy mode, core (or both edge and core) file services—2 GB RAM

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 16-9.

Configuring File Services

WAAS version 4.1.1 and later supports two mutually exclusive WAFS modes that are configured in different ways:

- Transparent CIFS accelerator mode—This mode requires no core, edge, or connectivity configuration. It relies on automatic discovery and transparently accelerates CIFS traffic with no configuration needed. We recommend this mode for ease of use.

  Table 11-2 provides an overview of the steps that you must complete to configure transparent CIFS accelerator mode. By using this mode, you can skip most of the configuration steps in this chapter.

- Legacy mode—This mode requires that you configure a Core cluster, Edge file services, and connections between the Core and Edge devices. Use this mode if you need to interoperate with WAAS 4.0.x devices or if disconnected mode operation is essential.

To switch an operating WAAS device between the two WAFS modes, see the "Switching File Services Modes" section on page 11-36.

Note Legacy mode WAFS is deprecated in WAAS version 4.2.1. It still operates but is not supported and will be removed in a future version. If you enable legacy mode WAFS, an alarm is raised on the device and you are alerted in all Central Manager GUI pages and the CLI if you configure any legacy WAFS settings. Legacy WAFS users should migrate to the transparent CIFS accelerator.

Table 11-3 provides an overview of the steps that you must complete to configure file services.

Table 11-2 Checklist for Configuring Transparent Mode File Services

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 11-8.
2. Enable CIFS acceleration.	Enables the transparent CIFS accelerator. For more information, see the "Enabling and Disabling the Global Optimization Features" section on page 12-2.
3. (Optional) Identify dynamic shares.	Identifies the dynamic shares on an exported file server. If your file server does not contain dynamic shares, you can skip this step. For more information, see the "Creating Dynamic Shares" section on page 11-21.
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 "Creating a Preposition Directive" section on page 11-25.

Table 11-3 Checklist for Configuring Legacy Mode File Services

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 11-8.
2. Configure a WAFS core cluster.	WAFS core clusters are required to copy data from the exported file server to the cache of the Edge WAEs. For more information, see the "Configuring the Core Cluster" section on page 11-11.
3. Configure the edge devices.	By default, file services are not enabled on your WAAS devices. To enable and start file services on edge devices, see the "Configuring the Edge Devices" section on page 11-14.
4. (Optional) Register a file server with the WAAS Central Manager.	Identifies to the WAAS system which file servers to export. This step also creates a link between a core cluster and the registered file server. For more information, see the "Setting Up File Servers to Export to the Edge WAE Cache" section on page 11-17. This step is optional if you are using automatic discovery.
5. (Optional) Identify dynamic shares.	Identifies the dynamic shares on a exported file server. If your file server does not contain dynamic shares, you can skip this step. For more information, see the "Creating Dynamic Shares" section on page 11-21.

Table 11-3 Checklist for Configuring Legacy Mode File Services (continued)

Task	Additional Information and Instructions
6. Create a connection between a core cluster and your edge devices.	Enables a core cluster to copy data to the Edge WAE cache. For more information, see the "Creating a Connection Between a Core Cluster and Edge WAEs" section on page 11-23.
7. (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 "Creating a Preposition Directive" section on page 11-25.

Configuring the Core Cluster

The first step in setting up file services in legacy mode is to enable Core services on a device and assign the device to a Core cluster. The Core cluster will be responsible for copying data from the exported file server (or multiple file servers) to the cache of the Edge WAEs. In later sections you will assign edge devices and, optionally, file servers to this core cluster, so the cluster knows which file servers to export and which edge devices to populate with cached data.

In Steps 1 through 7 you enable core server services and assign the device to a Core cluster. In Steps 8 through 12 you configure the new core cluster. The last step in this procedure describes how to reload the device, which is required for the device to function as a Core WAE.

To create a WAFS core cluster, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

The Devices window displays a list of devices created on your WAAS system.

Step 2 Click the **Edit** button next to the device that you want to be a member of the new core cluster.

The Device Dashboard window appears.

Step 3 From the navigation pane, choose **Configure > Acceleration > Legacy Services > WAFS Core Configuration**.

The Enable Core Server Services window appears. (See Figure 11-2.)

Note You cannot configure Core Server settings while CIFS accelerator is enabled.

Figure 11-2 Enabling Core Server Services

Step 4 Check the **Enable Core Server** check box.

Step 5 Assign this device to a core cluster by selecting the radio button next to that core cluster. If you do not have any existing core clusters, you must create a new legacy WAFS core cluster for this device. See the "Creating a New Device Group" section on page 3-3.

Step 6 Click **Submit**.

A popup message is displayed that the device must be manually rebooted for the device to function as a Core server.

Step 7 Click **OK** after reading the popup message.

Core Server services are enabled on the device and the device joins the specified Core cluster. The last step in this section describes how to reboot the device, which is required for the Core services to be activated.

If you click **Cancel** on the popup message you are returned to the Enable Core Server window and your changes are not submitted.

If you are relying on the automatic discovery feature instead of explicitly registering file servers, and you do not need to configure an optional differentiated services code point (DSCP) setting, you can skip to Step 14.

Step 8 Click the Home link at the top right in the WAAS Central Manager GUI to return to the global context, then from the navigation pane, choose **My WAN > Manage Device Groups**.

Step 9 Click the **Edit** icon next to the new core cluster.

Step 10 From the navigation pane, choose *Cluster Name* **> Core Server Settings**. The Configure Core Server Settings window appears. (See Figure 11-3.)

Figure 11-3 Configuring a Core Cluster Example

Step 11 In the File Server access username, File Server access password, and Confirm password fields, enter the access information that will be used for all of the CIFS file servers that are configured as part of this Core cluster. For prepositioning, the access credentials that you enter must allow read access to the prepositioned root directories and to their parent directories.

If you are relying on the automatic discovery feature instead of explicitly registering file servers, these fields are unnecessary.

You will specify which file servers this core cluster will export in the "Setting Up File Servers to Export to the Edge WAE Cache" section on page 11-17.

Step 12 (Optional) Configure a differentiated services code point (DSCP) value for high-priority messages, as follows:

a. Check the **Set DSCP value for high priority messages** check box.

b. Choose one of the following options:

- **Match Edge**—Matches the DSCP value of the Edge WAEs connected to this core cluster. This matching takes place when you create a connection between edge and core devices, as described in the "Creating a Connection Between a Core Cluster and Edge WAEs" section on page 11-23.
- **DSCP Value**—Allows you to specify a DSCP value for this core cluster.

  Select a value from the drop-down list and refer to Table 11-4 on page 11-16 for a description of the supported values. If you choose **Please Make a Choice** from the drop-down list, enter a value from 0 to 63 in the corresponding field.

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 (shown in Table 11-4) 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.

We recommend setting a DSCP value for WAFS control traffic to improve system performance. This requires your routers to be configured to enforce the QoS markings accordingly.

Step 13 Click **Submit**.

Step 14 Reload the device by clicking the **Reload WAE** icon in the taskbar, or by completing the following steps:

a. From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

b. Click the **Edit** icon next the device on which you enabled Core services. The Device Dashboard window appears.

c. Click the **Reload WAE** icon in the taskbar. The device is rebooted and Core services are activated on the device.

Configuring the Edge Devices

After you create and configure the core cluster, the next step (for legacy mode) is to configure the edge devices that will contain the exported file server data in their cache.

To enable the edge server on a device or device group, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Note We recommend enabling file services on an edge device group if WCCP is enabled on your network. If WCCP is disabled, you should enable file services on individual edge devices to prevent name conflicts.

Step 2 Click the **Edit** icon next to the edge device or device group on which you want to enable file services.

The Device Dashboard window or the Modifying Device Group window appears depending on the selected option.

You cannot enable edge services on the WAAS Central Manager device.

Step 3 From the navigation pane, choose **Configure > Acceleration > Legacy Services > WAFS Edge Configuration**.

The Enable Edge Server Services window appears. (See Figure 11-4.)

Note You cannot configure Edge Server Services settings while CIFS accelerator is enabled.

Figure 11-4 Enabling File Services on Edge Devices

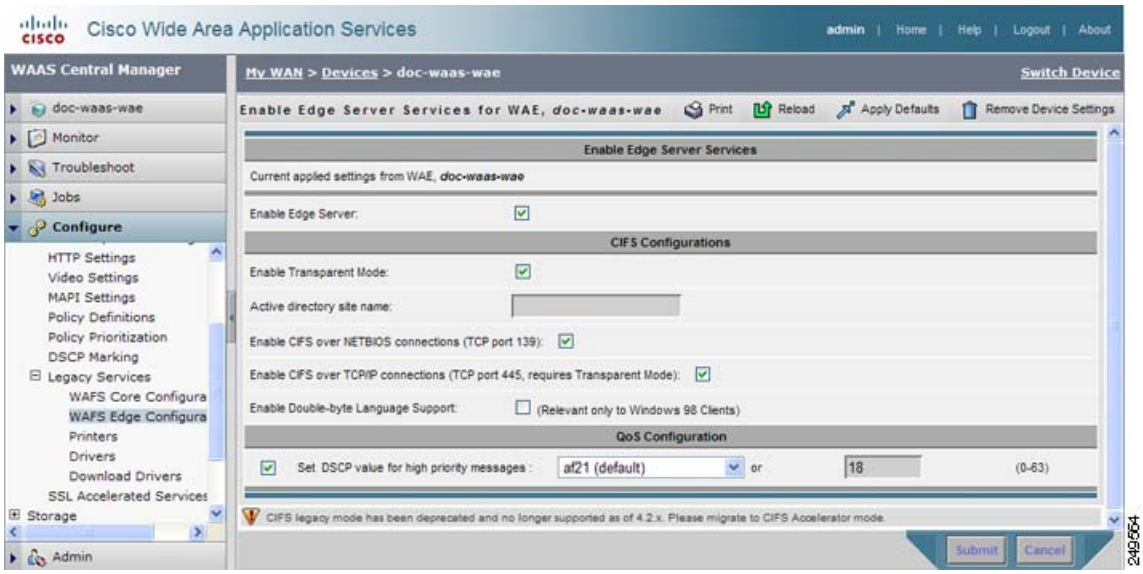

Step 4 Check the **Enable Edge Server** check box.

The other fields in the window are enabled.

Step 5 Check the **Enable Transparent Mode** check box if WCCP or PBR is enabled on your network, or if you are using inline mode.

The options for enabling TCP ports 139 and 445 are automatically updated based on whether transparent mode has been enabled.

If you enable transparent mode, the options for enabling TCP port 139 and port 445 are automatically selected. If you disable transparent mode (Enable Transparent Mode is not checked), the option for enabling TCP port 139 is selected and the option for enabling TCP port 445 is not selected, because port 445 is used only in transparent mode.

Step 6 Enter the name of the active directory site in the provided field.

Step 7 Enable the relevant ports on the Edge WAE by checking the following options (at least one must be checked):

- **Enable CIFS over NETBIOS connections (tcp port 139)**—Check this option if port 139 is open between your clients and the Edge WAEs, as well as between your core cluster and your file servers.

  If port 139 is not open on your network for security reasons, uncheck this option, and then complete the following tasks:

  – Enable WCCP on your routers and Edge WAEs, or enable inline mode on the Edge WAEs. For more information, see Chapter 4, “Configuring Traffic Interception.”

  – Enable port 445 on the Edge WAE by checking the **Enable CIFS Over TCP/IP Connections** check box.

- **Enable CIFS over TCP/IP connections (tcp port 445, requires Transparent Mode)**—Check this option if port 445 is open on your network**.**

  If port 445 is closed on your network, uncheck this option so that your Edge WAE does not try to establish a connection on this port, and then check the **Enable CIFS over NETBIOS connections** check box.

When you disable port 445 all clients connect directly to port 139 on the Edge WAE, then the core cluster connects to port 139 on the file server.

Note If you enable or disable connections on port 139 and port 445, existing clients will not lose their connection to the Edge WAEs.

Step 8 Check the **Enable Double-byte Language Support** check box if you have Windows 98 clients that need to support two-byte languages such as Japanese.

Leave this option unchecked in the following situations:

- You do not have any Windows 98 clients in your environment.
- You have Windows 98 clients in your environment, but they only need to support one-byte languages.

English will always be supported regardless of how this option is configured.

Step 9 (Optional) Configure a differentiated services code point (DSCP) value for high-priority messages, by following these steps:

a. Place a check in the **Set DSCP value for high priority messages** check box.

b. Select a value in the drop-down list. Refer to Table 11-4 for a description of the supported values.

If you choose **Please Make a Choice** from the drop-down list, enter a value from 0 to 63 in the corresponding field.

DSCP is a field in an IP packet that enables different levels of service to be assigned to network traffic. Each packet on the network is marked with a DSCP code (shown in Table 11-4) and is assigned to the corresponding level of service. DSCP is the combination of IP Precedence and Type of Service (ToS) fields. For more information, see RFC 2474.

We recommend setting a DSCP value for WAFS control traffic to improve system performance. This requires your routers to be configured to enforce the QoS markings accordingly.

Table 11-4 DSCP Codes

DSCP Code	Description
af11	Sets packets with AF11 dscp (001010).
af12	Sets packets with AF11 dscp (001100).
af13	Sets packets with AF13 dscp (001110).
af21	Sets packets with AF21 dscp (010010).
af22	Sets packets with AF22 dscp (010100).
af23	Sets packets with AF23 dscp (010110).
af31	Sets packets with AF31 dscp (011010).
af32	Sets packets with AF32 dscp (011100).
af33	Sets packets with AF33 dscp (011110).
af41	Sets packets with AF41 dscp (100010).
af42	Sets packets with AF42 dscp (100100).
af43	Sets packets with AF43 dscp (100110).
cs1	Sets packets with CS1 (precedence 1) dscp (001000).
cs2	Sets packets with CS2 (precedence 2) dscp (010000).
cs3	Sets packets with CS3 (precedence 3) dscp (011000).

Table 11-4 DSCP Codes (continued)

DSCP Code	Description
cs4	Sets packets with CS4 (precedence 4) dscp (100000).
cs5	Sets packets with CS5 (precedence 5) dscp (101000).
cs6	Sets packets with CS6 (precedence 6) dscp (110000).
cs7	Sets packets with CS7 (precedence 7) dscp (111000).
default	Sets packets with default dscp (000000).
ef	Sets packets with EF dscp (101110).

Step 10 Click **Submit**. A popup message is displayed that the device must be manually rebooted for the device to function as an Edge server.

Step 11 Click **OK** after reading the popup message. Edge Server services are enabled on the device.

If you click **Cancel** on the popup message, you are returned to the Edge Configuration window and your changes are not submitted.

Step 12 Reload the device by clicking the **Reload WAE** icon in the taskbar, or by completing the following steps:

a. From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

b. Click the **Edit** icon next the device or device group on which you enabled Edge services.

c. Click the **Reload WAE** or the **Reboot All Devices** icon in the taskbar. The device(s) are rebooted and Edge services are activated on the device(s).

Setting Up File Servers to Export to the Edge WAE Cache

For legacy mode, after you enable file services on a core cluster and Edge WAEs, you can optionally use the WAAS Central Manager GUI to set up the file servers that you want to export. If your network has many file servers that you need to define on the WAAS network (10 or more, for example), you can create and import a comma-separated values (CSV) file to speed up the process.

If you do not want to set up file servers but still want to use legacy mode, you can rely on the automatic discovery feature of legacy mode to allow WAFS to automatically discover file servers when they are accessed by users. If you are not explicitly registering file servers, you can skip this section.

This section contains the following topics on setting up a file server in the WAAS Central Manager GUI:

- Registering a File Server with the WAAS Central Manager, page 11-17
- Importing File Server Definitions Using a CSV File, page 11-18
- Assigning a Core Cluster to a Registered File Server, page 11-20
- Creating Dynamic Shares, page 11-21

Registering a File Server with the WAAS Central Manager

To register a file server with the WAAS Central Manager, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Configure > Legacy Services > File Servers**. The File Servers window appears.

From this window, you can perform the following tasks:

- Edit the configuration of an existing file server by clicking the **Edit** icon next to the file server. You can then delete the file server configuration, or modify any of the file server settings.
- Import multiple file server definitions using a CSV file by clicking the **Import from CSV** icon in the task bar. (See the "Importing File Server Definitions Using a CSV File" section on page 11-18.)
- Identify a new file server to export, as described in the next steps.

Step 2 Click the **Create New File Server** icon in the taskbar to identify a new file server to export. The Registering New File Server window appears.

Step 3 In the **File Server Name** field, enter the hostname of the file server to export. You can enter the server's NetBIOS name, DNS name, or IP address.

Note You cannot register the same file server more than once.

Step 4 Check the **Allow Access on WAN Failure** check box to provide CIFS clients with read-only access to the cached data on this Edge WAE in the event of a WAN failure.

When a WAN failure occurs, enabling this option allows CIFS clients to browse the cached directory structure and read fully cached files while authentication and authorization is maintained.

For more information, see the "Preparing Your WAAS Network for WAN Failures" section on page 11-34.

Step 5 Click **Submit**.

The file server is registered with the WAAS system, and the navigation pane refreshes with additional options.

Step 6 Proceed to the "Assigning a Core Cluster to a Registered File Server" section on page 11-20 to assign a core cluster to the registered file server.

Importing File Server Definitions Using a CSV File

If your network includes a high number (10 or more, for example) of file servers that you need to define on the WAAS network, you can create and import a comma-separated values (CSV) file to speed up the process. You can use Excel or another spreadsheet application to create the CSV file.

This section contains the following topics:

- CSV File Creation Requirements, page 11-18
- CSV File Considerations, page 11-19
- Importing the CSV File, page 11-19

CSV File Creation Requirements

Table 11-5 specifies the CSV file requirements.

Table 11-5 CSV File Requirements

Column Heading	Syntax/Semantics	Comments
Name	Specify the name of the file server. Corresponds to the **File Server Name** field in the Registering New File Server window.	Required. Same constraints as in Step 3 in Registering a File Server with the WAAS Central Manager.
AllowDisconnected	Corresponds to the Allow access on WAN Failure check box in the Registering New File Server window. Set to "true" to apply; set to "false" when not applicable.	Optional. If not specified, assumed to be "false."
Cluster	Specifies the name of an existing core cluster to which the file server is to be assigned.	Optional. If not specified, file server is not assigned to a core cluster. No error is reported if a cluster name is assigned to a file server more than one time.

CSV File Considerations

Keep the following in mind when creating CSV files:

- The first row must list the column headings that you are going to specify.
- At a minimum, the file must contain a "Name" column (other columns are optional).
- Column headings are not case sensitive and can be specified in any order.
- To assign a file server to multiple core clusters, you can specify multiple "Cluster" columns. The column headings row must contain multiple "Cluster" columns if you want to specify multiple clusters in the data rows.
- When you specify a row, it is not necessary to have a value for each column. However, you must specify the correct number of columns. That is, the number of columns must correspond to the number of objects defined in the headings row. The following example presents valid CSV file entries (assuming that core clusters c-1, c-2, and c-5 exist). The adjacent pairs of commas indicate that default values are to be used.

  Example:

  ```
  name,allowdisconnected,cluster,cluster
  s71,TRUE,c-1,
  s72,,c-2,c-5
  ```

Importing the CSV File

After you have created the CSV file, you use the WAAS Central Manager GUI to import the file.

To import the CSV file, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Configure > Legacy Services > File Servers**.

Step 2 In the taskbar, click the **Import from CSV** icon ( ).

The Importing File Server Definitions window appears.

Step 3 Click **Browse**.

The Choose File window opens.

Step 4 Navigate to and select the CSV file that you want to import, then click **Open**. The path and file server name appear in the Path to File Server Definitions field.

Step 5 In the Importing File Server Definitions window, click **Submit**.

The CSV file is imported and a "successfully imported" message is displayed.

The WAAS Central Manager checks the following:

- File headings are correct
- File server name appears on each row
- At least one file server is specified in the file
- All rows have the correct number of columns and each row is syntactically correct

If the file fails any of these checks, an error message appears in the WAAS Central Manager and no file servers are imported. Error messages may provide error explanations for up to 10 rows. However, if the error is related to the headings or the file cannot be read, no further checking occurs.

Assigning a Core Cluster to a Registered File Server

After you register a file server with the WAAS Central Manager, you need to assign at least one core cluster to the file server. The core cluster will be responsible for exporting the file server to the cache on the Edge WAEs. If you are relying on the automatic discovery feature rather than explicitly registering file servers, you do not need to perform this step.

To assign a core cluster to a registered file server, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Configure > Legacy Services > File Servers**.

The File Servers window appears.

Step 2 Click the **Edit** icon next to the file server that you want to assign to a core cluster.

Step 3 Click the **Core Clusters** tab. The Core Cluster Assignments window appears.

The Core Cluster Assignments window displays the device groups configured in your WAAS network. By default, the window displays 10 device groups.

Step 4 Assign a core cluster to the file server by doing either of the following:

- Click in the taskbar to assign all available core clusters to the file server.
- Click next to each core cluster that you want to assign to the file server. The icon changes to when selected.

Note You can only assign core clusters to file servers. You cannot assign regular device groups (identified by) to file servers.

Step 5 Click the **Resolve File Server Name** icon ( ) for the WAFS Core Cluster you selected. This icon is located next to the Type column and verifies that the name you entered for the file server can be resolved to an IP address.

This icon is available only for WAFS Core Clusters. If the file server name does not resolve, an error message appears in the Comments column. If this occurs, ensure you entered the correct name for the file server.

Step 6 Click **Submit**.

The icon next to the Core Clusters you selected changes to .

Step 7 (Optional) If the file server that you just added contains a dynamic share, see the "Creating Dynamic Shares" section on page 11-21.

When a file server contains a dynamic share, you must specify the dynamic share in the WAAS Central Manager GUI.

Creating Dynamic Shares

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, 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 Access Based Enumeration, if configured on 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.
- If you are using legacy mode, you can add two different dynamic shares with the same file server name and the same share name, but each one needs to be associated with a different core cluster. Each share will have a different ID.
- If you are using legacy mode, you can add dynamic shares only for explicitly registered file servers. Dynamic shares are supported on automatically discovered file servers only if you are using transparent CIFS accelerator mode.

To add a dynamic share, follow these steps:

Step 1 Before adding a dynamic share, verify the following (only if you are using legacy mode):

- The dynamic share is already set up on the CIFS file server.
- You have set up the file server in the WAAS Central Manager GUI. For more information on identifying a file server, see the "Setting Up File Servers to Export to the Edge WAE Cache" section on page 11-17.

Step 2 From the WAAS Central Manager GUI navigation pane, choose **Configure > 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 3 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 4 In the Name field, enter a name for the dynamic share.

The following characters are not supported in the dynamic share name: / \ : * ? " < > |
In legacy mode, the name can contain only letters, numbers, underscore (_), and hyphen (-).

Step 5 From the Assigned Domain drop-down list, choose the WAAS domain that you want to assign to the dynamic share. Only users who are also assigned to this domain have permission to access the dynamic share.

This kind of WAAS domain does not use entities. (When defining the domain, choose **None** for the Entity Type.) For more information on domains, see the "Working with Domains" section on page 7-14.

Step 6 (Optional) If you are using legacy mode, check the **CIFS - Use WAFS transport mode** check box.

Checking this box hides certain fields in this window because some fields are appropriate only for transparent CIFS accelerator mode.

Step 7 In the File Server field, enter the name of the file server with the dynamic share.

The text field is not shown if you checked the CIFS - Use WAFS transport mode check box. Instead, registered file servers are displayed in a drop-down list. For information on registering a file server, see the "Registering a File Server with the WAAS Central Manager" section on page 11-17. (Registering file servers is not required for transparent CIFS mode.)

Step 8 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 fields are not shown if you checked the CIFS - Use WAFS transport mode check box.)

These credentials are used only to access the file server when you click the **Browse** button.

Step 9 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 To make the Browse button appear in legacy mode, you must assign the file server to a Core Cluster and the cluster must contain at least one Core WAE. If these conditions are not met, the Browse button is not displayed in legacy mode. In transparent CIFS accelerator mode, 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 10 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 11 Click **Submit**.

The specified directory now functions as a dynamic share on the Edge WAE cache.

Creating a Connection Between a Core Cluster and Edge WAEs

After you have registered a file server with the WAAS system, you need to create a connection between a core cluster and your Edge WAEs (for legacy mode). This connection enables the core cluster to copy files to the cache on your Edge WAEs.

Before you define a connection that includes multiple core clusters and Edge WAEs, it is important to confirm that each core cluster-to-Edge WAE link has the same connection parameters, such as allocated bandwidth and roundtrip delay, as well as identical aliasing. If this is not the case, you must define a separate connection for each link.

To create a connection between a core cluster and one or more Edge WAEs, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Configure > Legacy Services > Connectivity**.

The Connectivity window appears.

From this window, you can perform the following tasks:

- Edit the configuration of an existing connection by clicking the **Edit** icon next to the connection. You can delete the connection, or modify any of the connection settings.
- Add a new connection, as described in the next steps.

Step 2 Click the **Create New Connection** icon in the taskbar to add a new connection.

The Creating a New Connection window appears.

Step 3 Enter a name for the connection.

Step 4 Choose the radio button next to the core cluster that you want included in this connection.

Step 5 Click **Submit**.

A message appears explaining that if you do not want to export any file servers for this connection, you need to uncheck each file server on the File Server Settings window before you assign an Edge device or group to this connection.

Step 6 Click **OK** after reading the message.

If you are relying on the automatic discovery feature instead of explicitly registering file servers, you do not need to configure file server settings and you can skip to Step 12.

Step 7 Click the **File Server Settings** tab.

The Configure File Server Aliases window appears. (See Figure 11-5.)

Figure 11-5 Configuring File Server Aliases

This window allows you to define the naming scheme for each file server. You can use the original file server name, the file server name plus a prefix or suffix, or an alias of your own design. For example, if you specify **as-** as a prefix and the name of the file server is **win3srv**, users will see this file server as **as-win3srv**.

This window also allows you to select the file servers that you want the core cluster to export.

Step 8 Specify a prefix or suffix value by selecting one of the following options:

- **Prefix**—Adds the prefix you enter to the beginning of the exported file server alias.
- **Suffix**—Adds the suffix you enter to the end of the exported file server alias.

The default behavior is for WAAS to add the prefix AS- to the beginning of the file server alias. You cannot enter a blank prefix or suffix value. If you specify an alias, as described in Step 10, it overrides the prefix and suffix setting.

Step 9 Check the check box next to each file server that you want to export in this connection (at least one must be selected). By default, all file servers are selected.

Note If no file servers appear in this window, you have not assigned this core cluster to a file server as described in the "Assigning a Core Cluster to a Registered File Server" section on page 11-20. You must complete the procedures in that section before proceeding.

Step 10 (Optional) Enter an alias for the selected file servers in the Alias column.

An alias, which can be any name (maximum of 15 characters), overrides the default prefix and suffix setting defined in Step 8.

Step 11 Click **Submit**.

The Exported As column displays the name of each exported file server as it will appear to your end users.

Step 12 Click one of the following tabs:

- **Assign Edge Devices** to assign individual edge devices to this connection.

- **Assign Edge Groups** to assign an edge device group to this connection.

The Edge Device Assignments window or the Edge Group 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 allows you to find items in the list that match the criteria that you set.

Step 13 Choose the edge devices to include in this connection by doing either of the following:

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

Note You can only assign edge devices or device groups with edge services enabled on them. You cannot assign regular core clusters or offline devices (identified by) to the connection. For information on enabling edge services, see the "Configuring the Edge Devices" section on page 11-14.

Step 14 Click **Submit**.

The icon next to each edge device or device group you selected changes to .

Step 15 (Optional) Configure WAN utilization settings for this connection, by completing the following tasks:

a. Click the **WAN Utilization** tab. The WAN Utilization window appears.

b. Define the following WAN utilization settings that control the bandwidth that is used for WAFS traffic between the Edge and Core WAEs:

- **Maximum allocated bandwidth**—Enter the maximum bandwidth allocated to the connection (in kilobits per second). This value must be less than or equal to the maximum bandwidth of the physical WAN link between the WAEs in this connection. This setting has a default value of 1544 KB per second. WAFS throttles its bandwidth usage to 1.5 times the value of this setting.
- **Minimum roundtrip delay**—Enter the length of time (in milliseconds) it takes a bit to travel round-trip from one WAE to the other end and back when the link is idle. This setting has a default value of 80 msec.

Note If you later change any of the WAN utilization settings for this connection, you need to restart the Edge WAE for the new values to take effect.

Step 16 Click **Submit**.

Creating a Preposition Directive

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.

You can add preposition directives only for explicitly registered file servers in legacy mode. Prepositioning is supported on automatically discovered file servers only in transparent CIFS accelerator mode.

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. For transparent CIFS accelerator mode, these items are specified directly in the Creating New Preposition Directive window, as described in the following procedure. For details on specifying the username and password for legacy mode, see Step 11 in the "Configuring the Core Cluster" section on page 11-11.

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 11-27.

When using prepositioning in transparent CIFS accelerator mode, both branch and data center WAEs are required (the same as for any other accelerated traffic). 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 A warning message appears if the required connections do not exist for defining preposition directives in legacy mode.

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 a preposition directive:

- Creating a New Preposition Directive, page 11-27
- Assigning Edge Devices to a Preposition Directive, page 11-31

- Creating a New Preposition Schedule, page 11-32

Creating a New Preposition Directive

To create a preposition directive, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Configure > 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 11-6.)

Figure 11-6 Creating a New Preposition Directive

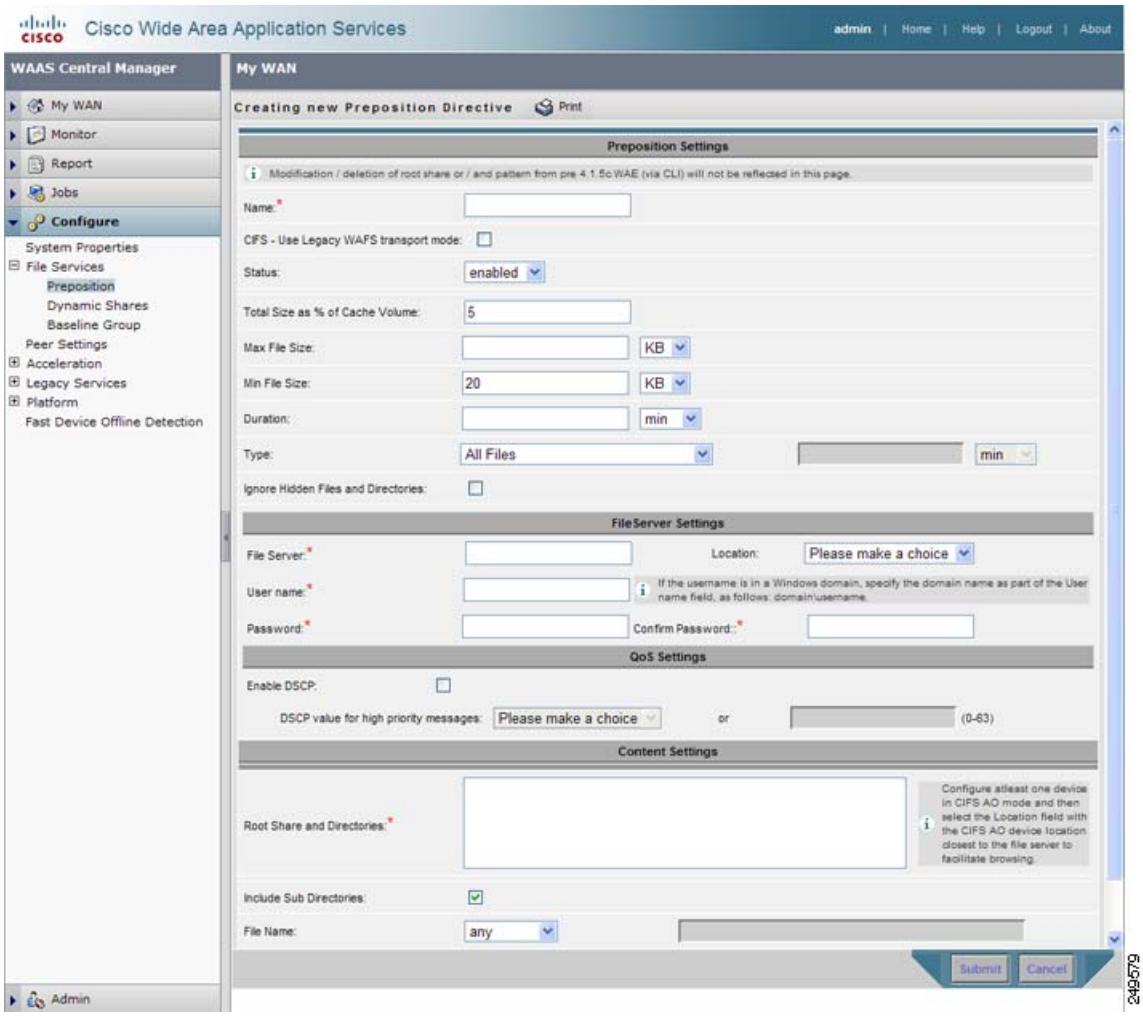

Step 3 Enter a name for the directive. The double quote (") character is not allowed in the name.

Step 4 (Optional) If you are using legacy mode, check the **CIFS - Use Legacy WAFS transport mode** check box.

Checking this box hides certain fields in this window because some fields are appropriate only for transparent CIFS accelerator mode.

Step 5 From the Status drop-down list, choose either **enabled** or **disabled**. Disabled directives are not put into effect.

Step 6 (Optional) In the File Server drop-down list, choose the name of a file server to export. For information on registering a file server, see the "Registering a File Server with the WAAS Central Manager" section on page 11-17. Registering file servers is not required for transparent CIFS mode. (The File Server field is not shown if you checked the CIFS - Use Legacy WAFS transport mode check box.)

Step 7 (Optional) Define the time and size limitations using the provided fields.

Table 11-6 describes the time and size limitation fields.

Table 11-6 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 8 (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.

Note If you are using legacy mode and relying on the automatic discovery feature instead of explicitly registering file servers, this check box does not apply because prepositioning is not supported for automatically discovered file servers in legacy mode.

Step 9 In the File Server field, enter the name of a file server to export. Do not use the double quote (") or forward slash (/) characters.

The File Server Settings area is not shown if you checked the CIFS - Use Legacy WAFS transport mode check box. Instead, registered file servers are displayed in a drop-down list (see Step 6).

Step 10 From the Location drop-down list, choose the device location that will provide browsing services for the file server. For the best browsing performance, specify a location that is close to the file server. For more information on defining locations, see the "Working with Device Locations" section on page 3-14. (The Location is not shown if you checked the CIFS - Use Legacy WAFS transport mode check box.)

Step 11 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 fields are not shown if you checked the CIFS - Use Legacy WAFS transport mode check box.)

The access credentials that you enter must allow read access to the prepositioned root directories and to their parent directories.

Step 12 (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. See Table 11-4 on page 11-16 for a description of the supported values. (The DSCP settings are not shown if you checked the CIFS - Use Legacy WAFS transport mode check box.)

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 13 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.

  In legacy mode, the Browse button appears only if you have configured the File Server access username and the File Server access password fields in the Core Server Settings configuration window. (See Figure 11-3.) In transparent CIFS mode, 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 14 Click **Submit**.

The directive is saved to the system and additional options appear in the navigation pane.

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 11-27.

To assign an Edge WAE or device group to a preposition directive, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Configure > 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 mode preposition directive to a device group, the directive is applied only to those devices enabled for CIFS acceleration in the assigned device group. Similarly, when assigning a legacy mode prepostion directive to a device group, the directive is applied only to those devices enabled for legacy mode 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 GUI navigation pane, choose **Configure > 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.

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 11-33.

Managing File Services

The following topics in this section describe how to manage file servers:

- Checking the Preposition Status, page 11-33
- Starting and Stopping Preposition Tasks, page 11-34
- Preparing Your WAAS Network for WAN Failures, page 11-34
- Viewing Members of a Core Cluster, page 11-36
- Switching File Services Modes, page 11-36

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 GUI navigation pane, choose **Configure > 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 10-20.

Preparing Your WAAS Network for WAN Failures

When you set up a connectivity directive that links an Edge WAE to a core cluster, you have the option to configure the Edge WAE to operate in disconnected mode in the event a WAN failure breaks its link to the core cluster.

Note Disconnected mode is not available when you are using the transparent CIFS accelerator mode of file services.

This section contains the following topics:

- About Disconnected Mode, page 11-35
- DNS and Domain Controller Requirements, page 11-35
- Data Availability in Disconnected Mode, page 11-36
- Configuring Disconnected Mode, page 11-36

About Disconnected Mode

Disconnected mode allows CIFS clients to continue to browse the cache directory and read fully cached files on an Edge WAE when a WAN failure occurs. Because the Edge WAE cannot verify its cached data against the file server during a WAN failure, CIFS clients are provided with read-only access to the cached data.

When the WAN connection between the Edge WAE and core cluster is restored, the Edge WAE automatically switches back to regular connected mode.

An Edge WAE switches to disconnected mode when the Edge WAE loses its connection to the core cluster. This is known as a WAN failure.

Note A file server crash does not trigger a switch to disconnected mode (assuming the Edge WAE maintains its connection to the core cluster).

An Edge WAE can operate in disconnected mode for one connection and in regular connected mode for another connection. For example, if you create two connectivity directives for an Edge WAE where one directive links it to core cluster A and the other directive links it to core cluster B and a WAN failure breaks its link to core cluster A, then the Edge WAE switches to disconnected mode for that connection. The Edge WAE will remain in regular connected mode when interacting with core cluster B.

Disconnected mode operates only with explicitly registered file servers. Disconnected mode is not supported for automatically discovered file servers.

DNS and Domain Controller Requirements

Disconnected mode requires a local domain controller in the branch office to authenticate CIFS clients. An Edge WAE operating in disconnected mode supports all authentication methods except Kerberos.

In DNS-only environments, a local DNS server at the branch office is also required. If WINS is used instead of DNS, there must be a local WINS server in the branch office.

To enable disconnected mode access to cached files, you must add the Edge WAE to the Active Directory or Windows NT domain (to perform windows authentication in case of WAN failure).

Note By default, Windows domain controllers enforce automatic machine account password changes as part of the authentication process. The machine account password for the Edge WAE is automatically negotiated and changed between the Edge WAE and the domain controller every seven days. However, if the authentication service is down, this process may not occur, and the machine account password for the Edge WAE may expire. To prevent this situation, we recommend that you disable automatic machine account password changes for the Edge WAE. For more information, see the "Disabling the Automatic Machine Account Password Changes for the Edge WAE" section on page 6-23.

Note WINS registrations usually have a three-day timeout. As a result, if WCCP is enabled in disconnected mode there is a chance that the WINS registration of the original file servers will expire and clients will experience name resolution problems.

Data Availability in Disconnected Mode

When an Edge WAE is in disconnected mode, CIFS clients are allowed read-only access to fully cached files on the Edge WAE. CIFS clients are able to view partially cached files and non-cached files in the directory structure, but they are not able to open these files.

Directories are available in disconnected mode only if their ACL is cached. This means that shares that have never been accessed, and therefore are not in the cache, are not displayed.

The best way to ensure a file is available in disconnected mode is to set up a preposition directive that proactively places a copy of the file on the Edge WAE cache. For more information, see the "Creating a New Preposition Directive" section on page 11-27.

Configuring Disconnected Mode

When you register a file server with the WAAS Central Manager, you can configure an Edge WAE to operate in disconnected mode. For more information, see the "Registering a File Server with the WAAS Central Manager" section on page 11-17.

You also must add the Edge WAE to the Active Directory or Windows NT domain. For details, see the "Configuring Windows Domain Server Settings on a WAAS Device" section on page 6-18. Verify that the WAE was successfully added into the domain by clicking the **Show Authentication Status** button. Even if the authentication status is not OK, as long as the **wbinfo -t** command executed successfully (trust via RPC is successful), the WAE will provide authentication for Read Only disconnected mode.

Viewing Members of a Core Cluster

Note This functionality is not available when you are using the transparent CIFS accelerator mode of file services because Core Clusters do not exist in this mode.

To view the Core WAEs that are members of a Core Cluster device group, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Device Groups**. The Device Groups window appears.

Step 2 Click the **Edit** icon next to the WAFS Core Cluster device group for which you want to view its members. The Modifying Device Group window appears.

Step 3 From the navigation pane, choose **Members**. A list of the devices that belong to the selected device group appears.

Switching File Services Modes

WAAS version 4.1.1 and later supports two mutually exclusive WAFS modes that are configured in different ways: transparent CIFS accelerator mode and legacy mode. This section describes how to switch modes on a WAAS device and contains the following topics:

- Changing from Legacy Mode to Transparent CIFS Accelerator Mode, page 11-37
- Changing from Transparent CIFS Accelerator Mode to Legacy Mode, page 11-38

For additional details about migrating from legacy mode to CIFS transparent accelerator mode, see the document *Cisco Wide Area Application Services Software 4.1 Common Internet File System Migration*.

Note Legacy mode WAFS is deprecated in WAAS version 4.2.1. It still operates but is not supported and will be removed in a future version. If you enable legacy mode WAFS, an alarm is raised on the device and you are alerted in all Central Manager GUI pages and the CLI if you configure any legacy WAFS settings. Legacy WAFS users should migrate to the transparent CIFS accelerator.

Changing from Legacy Mode to Transparent CIFS Accelerator Mode

If the WAAS device is currently using legacy mode and you want to change it to transparent CIFS accelerator mode, follow these steps:

Step 1 If there are existing preposition directives, edit the preposition directives to remove the assigned edge devices:

a. From the WAAS Central Manager GUI navigation pane, choose **Configure > File Services > Preposition**. The Preposition Directives window appears.

b. Click the **Edit** icon next to the preposition directive that you want to edit.

c. Click the **Assign Edge Devices** tab to display the assigned edge devices.

d. Click the **Remove all Edge Devices** icon to remove all edge devices.

e. Click the Assign Edge Groups tab to display the assigned edge device groups.

f. Click the **Remove all Edge Groups** icon to remove all edge device groups.

Step 2 If there are existing preposition directives, edit the preposition directives to change their mode to transparent CIFS accelerator mode:

a. Uncheck the box **CIFS - Use WAFS transport mode** and enter the file server name in the File Server field (for details on editing a preposition directive, see the "Creating a New Preposition Directive" section on page 11-27).

b. From the Location drop-down list, choose the device location that will provide browsing services for the file server. For the best browsing performance, specify a location that is close to the file server. For more information on defining locations, see the "Working with Device Locations" section on page 3-14.

c. In the User name, Password, and Confirm 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.

d. Click **Submit**.

Step 3 If there are existing dynamic shares defined, edit the dynamic shares to change their mode to transparent CIFS accelerator mode (for details on editing a dynamic share, see the "Creating Dynamic Shares" section on page 11-21):

a. Uncheck the box **CIFS - Use WAFS transport mode** and enter the file server name in the File Server field.

b. In the User name, Password, and Confirm 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.

c. Click **Submit**.

Step 4 Delete the connectivity directives that have been defined between core clusters and edge WAEs (see the "Creating a Connection Between a Core Cluster and Edge WAEs" section on page 11-23 for more information):

a. From the WAAS Central Manager GUI navigation pane, choose **Configure > Legacy Services > Connectivity**.

b. Click the **Edit** icon next to each connectivity directive and then click the **Delete Connectivity Directive** icon in the taskbar. Click **OK** in the confirmation dialog.

Step 5 Disable Edge Server Services and Core Server Services on the device, as described in the "Configuring the Edge Devices" section on page 11-14 and in the "Configuring the Core Cluster" section on page 11-11. You must uncheck the Enable Edge Server and Enable Core Server check boxes.

Step 6 Enable the transparent CIFS accelerator:

a. From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

b. Click the **Edit** icon next to the device or device group for which you want to change the global optimization features.

c. From the navigation pane, choose **Configure > Acceleration > Enabled Features** (see the "Enabling and Disabling the Global Optimization Features" section on page 12-2).

d. Check CIFS Accelerator.

e. Click **Submit**.

Step 7 Reload the device and wait for it to restart.

Step 8 Assign edge devices to any existing or new preposition directives, as described in the "Assigning Edge Devices to a Preposition Directive" section on page 11-31.

Note If you switch file services modes by using the CLI, you may see the following false alarm in the Central Manager GUI: Roles changed, reload required. To avoid this alarm, wait two data feed poll cycles (about 10 minutes) before reloading the device after switching modes from the CLI.

Changing from Transparent CIFS Accelerator Mode to Legacy Mode

If the WAAS device is currently using transparent CIFS accelerator mode and you want to change it to legacy mode, follow these steps:

Step 1 Define a file server as described in the "Registering a File Server with the WAAS Central Manager" section on page 11-17.

Step 2 If there are existing preposition directives, edit the preposition directives to remove the assigned edge devices:

a. From the WAAS Central Manager GUI navigation pane, choose **Configure > File Services > Preposition**. The Preposition Directives window appears.

b. Click the **Edit** icon next to the preposition directive that you want to edit.

c. Click the **Assign Edge Devices** tab to display the assigned edge devices.

d. Click the **Remove all Edge Devices** icon to remove all edge devices.

e. Click the Assign Edge Groups tab to display the assigned edge device groups.

f. Click the **Remove all Edge Groups** icon to remove all edge device groups.

Step 3 If there are existing preposition directives, edit the preposition directives to change their mode to legacy mode by checking the box **CIFS - Use WAFS transport mode** and selecting a file server from the File Server drop-down list. Click **Submit**. For details on editing a preposition directive, see the "Creating a New Preposition Directive" section on page 11-27.

Step 4 If there are existing dynamic shares defined, edit the dynamic shares to change their mode to legacy mode by checking the box **CIFS - Use WAFS transport mode** and selecting a file server from the File Server drop-down list. Click **Submit**. For details on editing a dynamic share, see the "Creating Dynamic Shares" section on page 11-21.

Step 5 Disable the transparent CIFS accelerator:

a. From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

b. Click the **Edit** icon next to the device or device group for which you want to change the global optimization features.

c. From the navigation pane, choose **Configure > Acceleration > Enabled Features** (see the "Enabling and Disabling the Global Optimization Features" section on page 12-2).

d. Uncheck CIFS Accelerator.

e. Click **Submit**.

Step 6 Enable Edge Server Services on the device, as described in the "Configuring the Edge Devices" section on page 11-14, or, if the device is to be used as a core device, enable Core Server Services on the device, as described in the "Configuring the Core Cluster" section on page 11-11.

Step 7 Reload the device and wait for it to restart.

Step 8 Assign a Core Cluster to each registered file server, as described in the "Assigning a Core Cluster to a Registered File Server" section on page 11-20. This step assumes that you already have created a Core Cluster.

Step 9 Create connectivity between the edge and core devices, as described in the "Creating a Connection Between a Core Cluster and Edge WAEs" section on page 11-23.

Step 10 Assign edge devices to any preposition directives, as described in "Assigning Edge Devices to a Preposition Directive" section on page 11-31.

CHAPTER 12

Configuring Application Acceleration

This chapter describes how to configure the application 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, page 12-1
- Enabling and Disabling the Global Optimization Features, page 12-2
- Creating a New Traffic Application Policy, page 12-29
- Managing Application Acceleration, page 12-37

About Application Acceleration

The WAAS software comes with over 150 predefined application 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 Application Policies."

Each application policy contains the following elements:

- Application definition—Identifies general information about a specific application, such as the application name, the differentiated services code point (DSCP) marking value that is applied to the traffic, and whether the WAAS Central Manager collects statistics about this application.
- Classifier—Contains a matching condition that identifies specific types of traffic. For example, the default HTTP classifier matches all traffic going to ports 80, 8080, 8000, 8001, and 3128. You can create up to 512 classifiers and 1024 matching conditions.
- Policy—Combines the application definition and classifier 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 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 application policies, see the "Creating a New Traffic Application Policy" section on page 12-29. For more information on viewing reports, restoring policies, monitoring applications, and other functions, see the "Managing Application Acceleration" section on page 12-37.

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.

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, and video. By default, all of the application accelerators are enabled. The application accelerators also require specific licenses to operate. For information on installing licenses, see the "Managing Software Licenses" section on page 9-3.

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group for which you want to change the global optimization features.

Step 3 From the navigation pane, choose **Configure > Acceleration > Enabled Features**.

The Enabled Features window is displayed. (See Figure 12-1.)

Figure 12-1 Modifying the Global Optimization Features

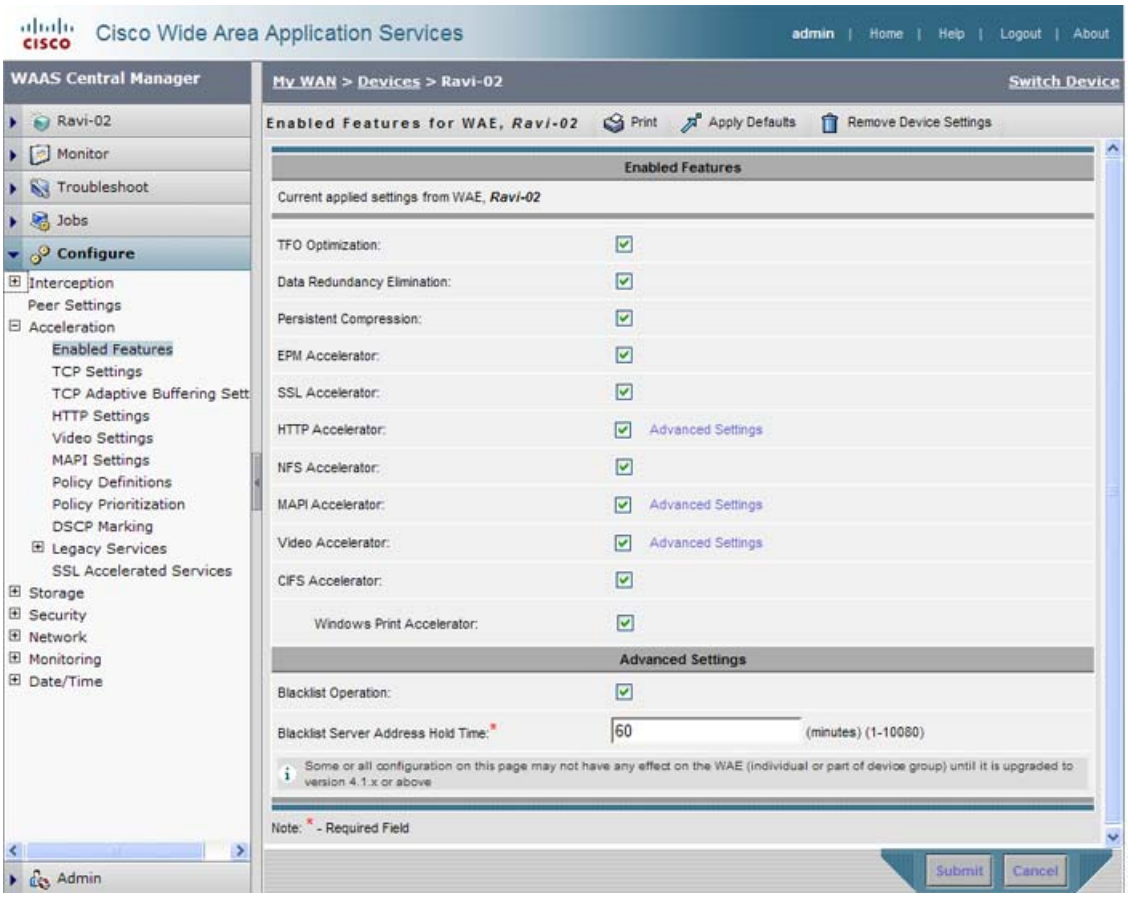

Step 4 Place a check next to the optimization features that you want to enable, and uncheck the features that you want to disable. For a description of each of the optimization features, see the "Key Services of Cisco WAAS" section on page 1-4.

Step 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 12-5.

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 12-9.

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 12-8.

Step 8 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 If you are changing from WAFS legacy mode to the CIFS accelerator, you must disable the WAFS legacy mode before you can enable the CIFS accelerator. To disable WAFS legacy mode, you must disable Data Center and Branch file services. For WAFS configuration information, see Chapter 11, "Configuring Wide Area File Services."

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 9 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 10 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 11 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 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 using individual application accelerators, see the following sections:

- Configuring HTTP Acceleration, page 12-5
- Configuring MAPI Acceleration, page 12-8
- Configuring Video Acceleration, page 12-9
- Configuring SSL Acceleration, page 12-11
- For CIFS: Chapter 11, "Configuring Wide Area File Services"

Configuring HTTP Acceleration

The HTTP application accelerator accelerates HTTP traffic. SSL traffic that uses HTTPS is handed off to the SSL accelerator.

The default Web application policy is defined to send traffic to the HTTP accelerator. The Web application policy uses the HTTP classifier, 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 classifier.

To enable the HTTP accelerator, check the HTTP Accelerator check box in the Enabled Features window (see Figure 12-1).

To configure the HTTP acceleration settings, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group for which you want to change the HTTP acceleration configuration.

Step 3 From the navigation pane, choose **Configure > Acceleration > HTTP Settings**.

The HTTP Acceleration Settings window is displayed. (See Figure 12-2.)

Figure 12-2 Modifying the HTTP Acceleration Settings

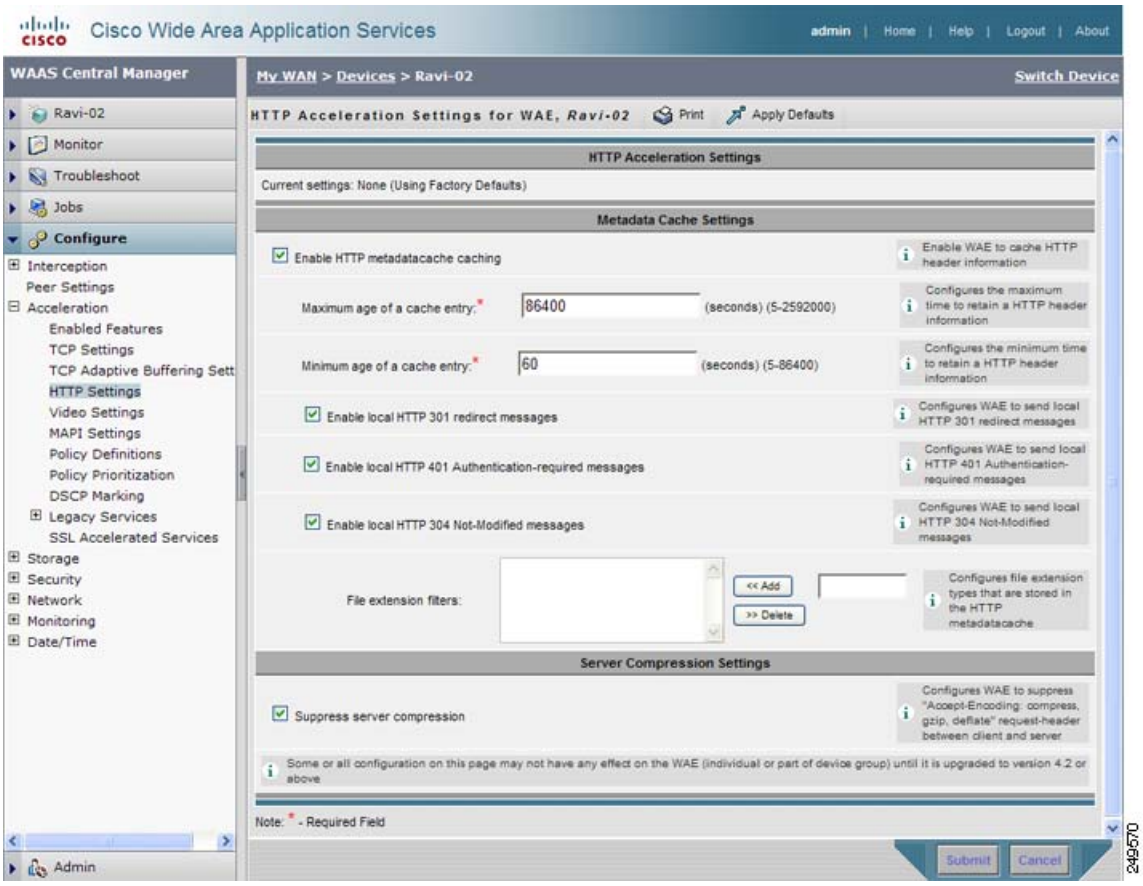

Step 4 Check the Enable HTTP metadatacache caching check box to enable the WAE to cache HTTP header (metadata) information. The default setting is unchecked.

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 12-7.

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 check box to configure the WAE to suppress server compression between the client and the server. The default setting is unchecked.

By checking this box, you are telling the WAE to remove the Accept-Encoding value from HTTP request headers, preventing the web server from compressing HTTP data that it sends to the client. This allows the WAE to apply its own compression to the HTTP 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 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.

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.

Configuring MAPI Acceleration

The MAPI application accelerator accelerates Microsoft Outlook Exchange traffic that uses the Messaging Application Programming Interface (MAPI) protocol. Microsoft Outlook 2000–2007 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) or encryption are not accelerated, and MAPI over HTTP is not accelerated.

Note Microsoft Outlook 2007 has 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 application 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 application policy that accelerates MAPI traffic with a classifier 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 12-1).

To configure MAPI acceleration settings, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group for which you want to change the MAPI acceleration configuration.

Step 3 From the navigation pane, choose **Configure > Acceleration > MAPI Settings**.

The MAPI Acceleration Settings window is displayed. (See Figure 12-3.)

Figure 12-3 Modifying MAPI Acceleration Settings

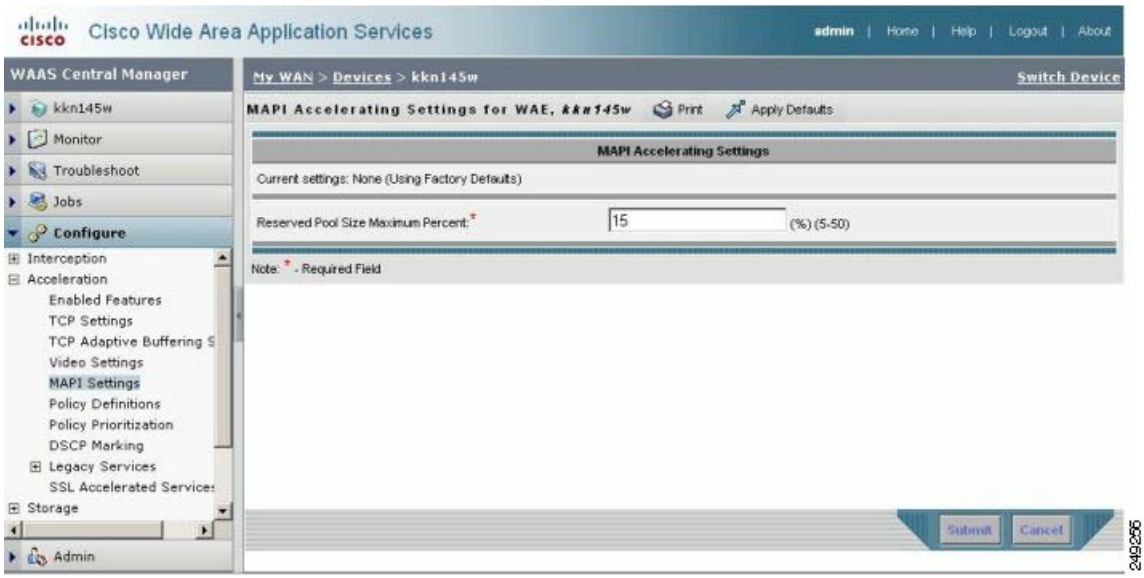

Step 4 In the **Reserved Pool Size Maximum Percent** field, enter the maximum percent of connections to restrict the maximum number of connections reserved for MAPI optimization during TFO overload. It is specified as a percent of the TFO connection limit of the platform. Valid percent ranges from 5%-50%. 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 5 Click **Submit**. The changes are saved to the device or device group.

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 classifier for the Streaming application 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 12-1).

To configure the video acceleration settings, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group for which you want to change the video acceleration configuration.

Step 3 From the navigation pane, choose **Configure > Acceleration > Video Settings**.

The Video Acceleration Configuration window is displayed. (See Figure 12-4.)

Figure 12-4 Modifying the Video Acceleration Configuration

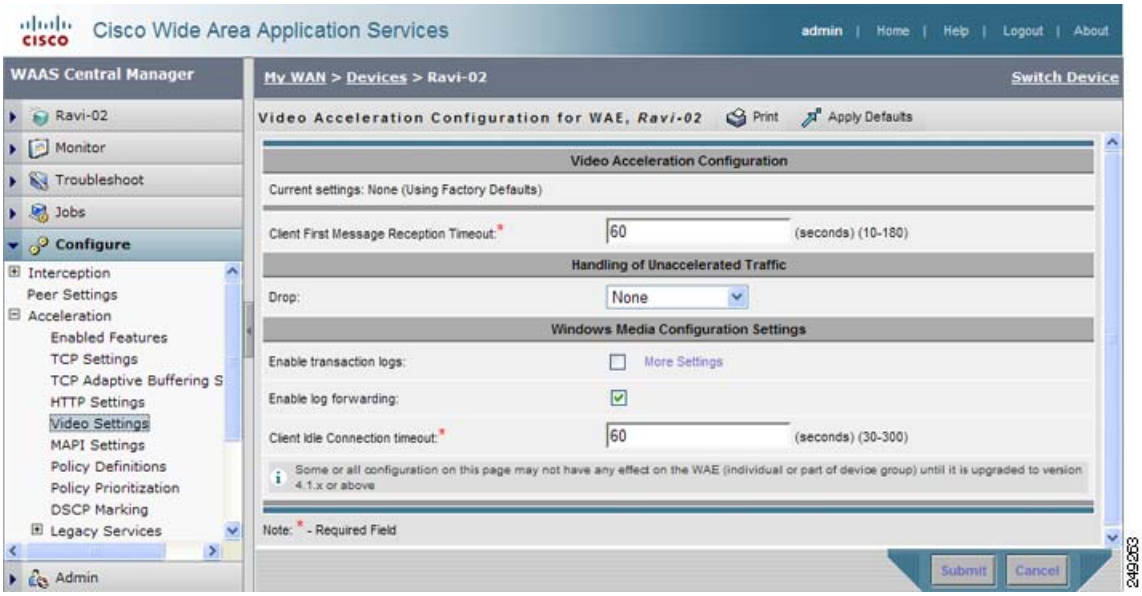

Step 4 In the Client First Message Reception Timeout field, enter the number of seconds to wait for the first message from the client and the first response from the server, after the connection is accepted by the video accelerator, before timing out the connection. Valid values range from 10–180 seconds. The default is 60.

Step 5 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 6 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 7 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 8 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 9 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 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 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, and then hands the connection to SSL accelerator for optimization).

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 12-1 provides an overview of the steps you must complete to set up and enable SSL acceleration.

Table 12-1 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 12-12.

Table 12-1 Checklist for Configuring SSL Acceleration (continued)

Task	Additional Information and Instructions
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 12-13.
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 12-2.
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 12-14.
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 12-19.
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 12-21.
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 12-23.
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 12-25.
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 12-26.

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 12-5 shows how the WAAS software handles SSL application optimization.

Figure 12-5 SSL Acceleration Block Diagram

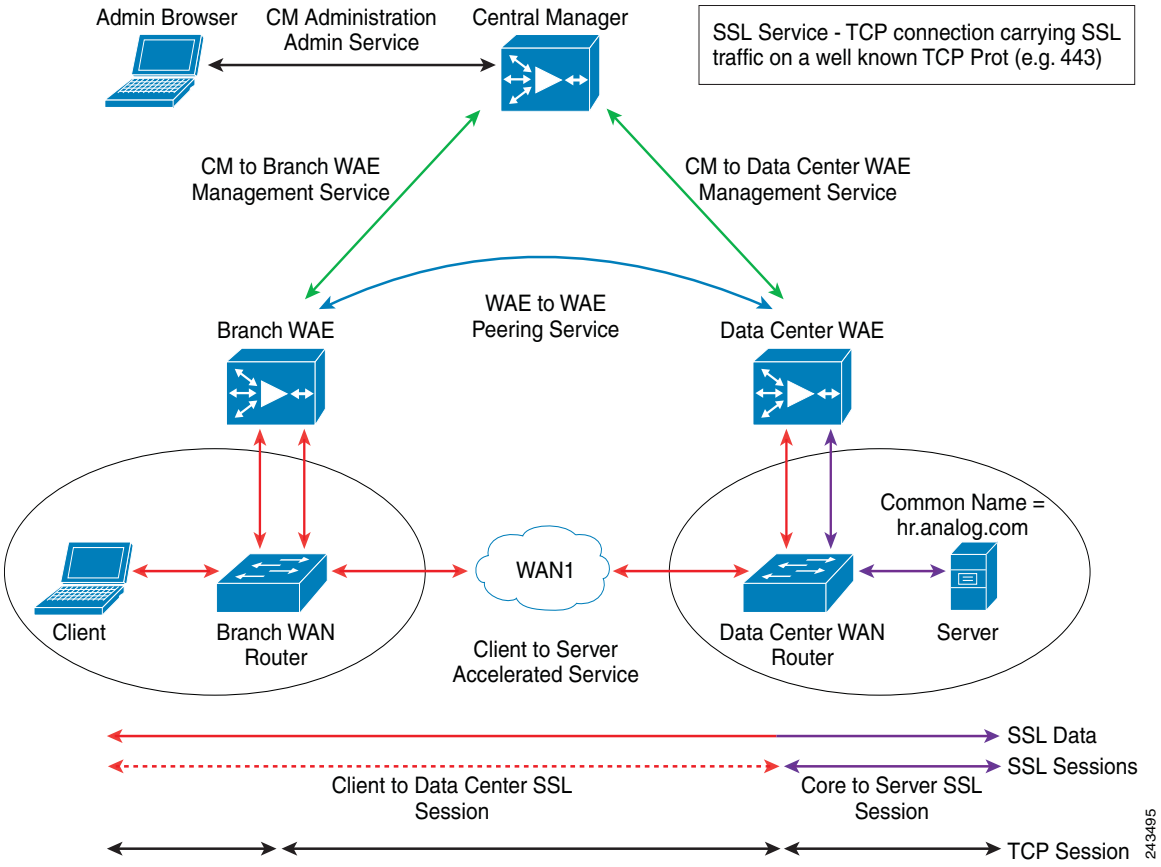

When you configure SSL acceleration, you must configure SSL accelerated service on the server-side (Data Center) WAE devices. The client-side (Branch) WAE needs to have its secure store initialized and unlocked/opened, but does not need to have the SSL accelerated service configured. However, the SSL 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 9-10.

Step 2 Enable the Enterprise license.

To enable the Enterprise license, see the "Managing Software Licenses" section on page 9-3.

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 12-2.

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group for which you want to configure SSL acceleration.

Step 3 From the navigation pane, choose **Configure > Security > SSL > Global Settings**.

The SSL Global Settings window is displayed (see Figure 12-6).

Figure 12-6 SSL Global Settings

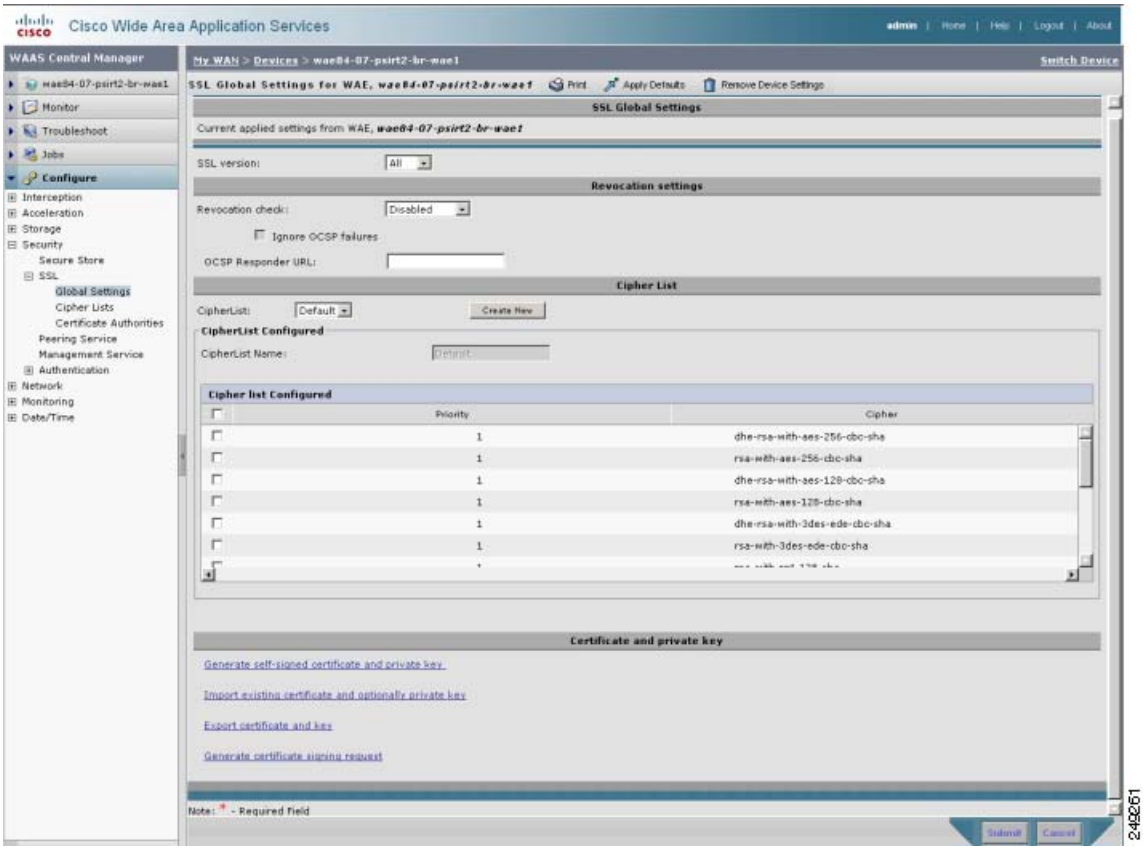

Step 4 To configure a device to use the SSL settings from a particular device group, choose the device group from **Select a Device Group** drop-down list located in SSL global settings toolbar. A device can either use its own SSL settings, or SSL settings from a device group. However, it is not possible to configure a device to use SSL settings from multiple device groups.

Step 5 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 6 (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 7 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 12-19.

Step 8 Choose a certificate/key pair method (see Figure 12-7).

Figure 12-7 Configuring Service Certificate and Private Key

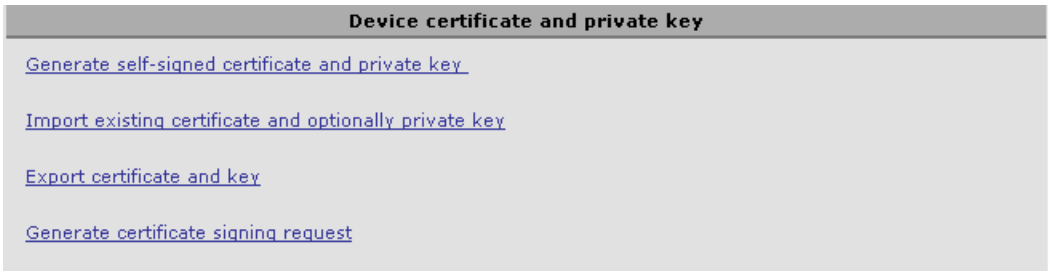

- Click **Generate Self-signed Certificate Key** to have the WAAS devices use a self-signed certificate/key pair for SSL.
- Click **Import Existing Certificate Key** to upload or paste an existing certificate/key pair.
- Click **Export Certificate Key** to export the current certificate/key pair.
- Click **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 12-15.

Step 9 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 12-8), follow these steps:

Figure 12-8 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 12-9), follow these steps:

Note WAAS SSL feature only supports RSA signing/encryption algorithm and keys.

Figure 12-9 Importing Existing Certificate or Certificate Chain

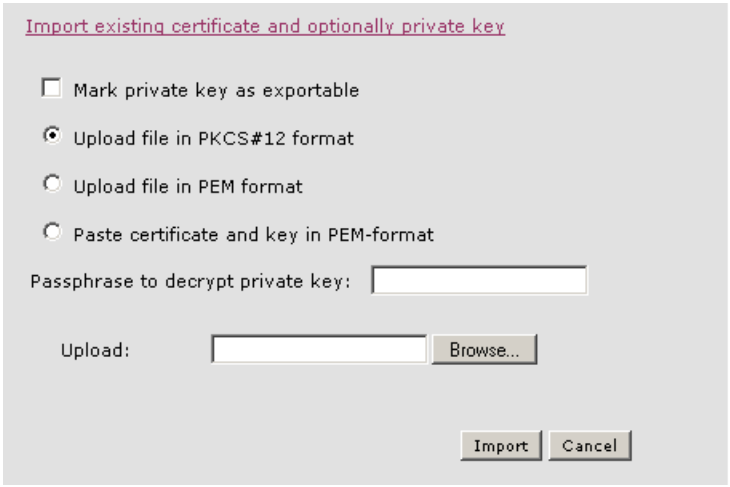

a. Check the **Mark private key as exportable** check box to export this certificate/key in the WAAS Central Manager and device CLI later.

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 12-10), follow these steps:

Figure 12-10 Export Certificate and Key

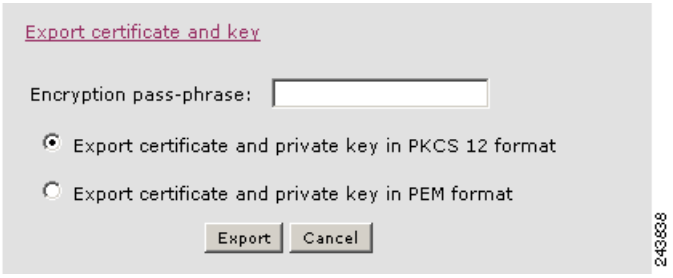

a. Enter the encryption pass-phrase.

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 12-11), follow these steps:

Figure 12-11 Generate Certificate Signing Request

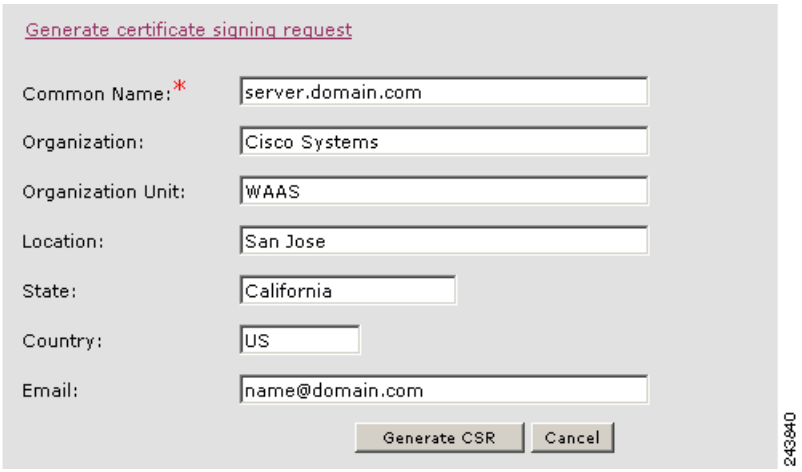

To update the current certificate with one signed by the Certificate Authority:

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group for which you want to configure cipher lists.

Step 3 From the navigation pane, choose **Configure > Security > SSL > Cipher Lists**.

The SSL Cipher Lists window is displayed (see Figure 12-12).

Figure 12-12 Displaying the SSL Cipher Lists

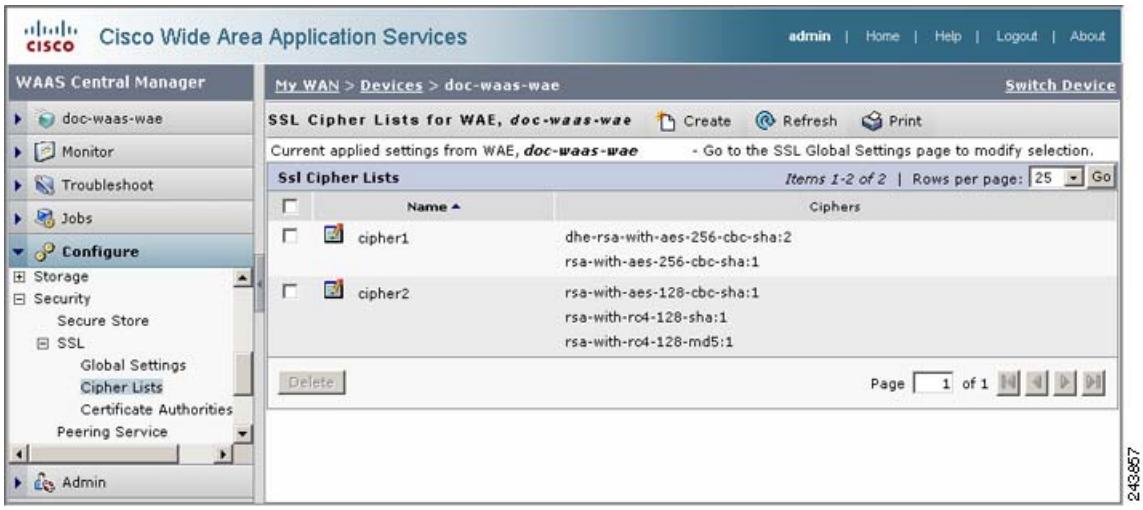

Step 4 Click **Create** to add a new cipher list.

The Creating New SSL Cipher List window is displayed (see Figure 12-13).

Figure 12-13 Creating a SSL Cipher List

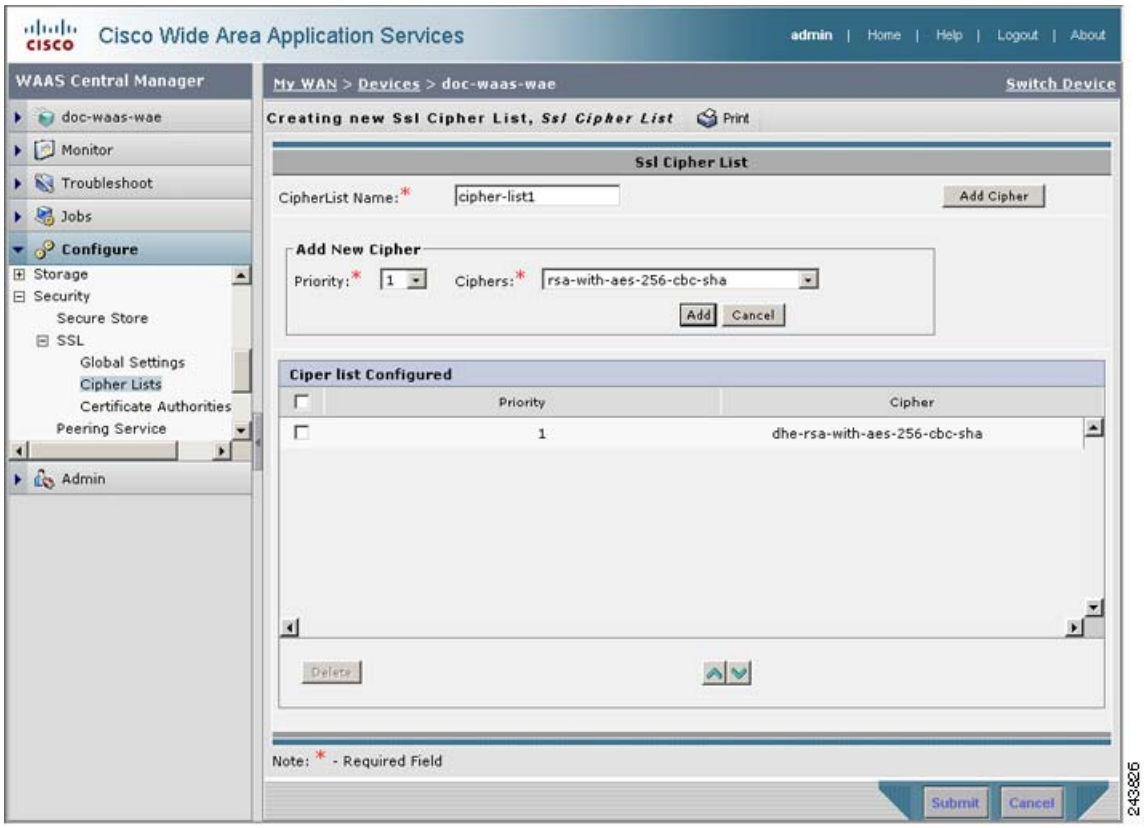

Step 5 Type a name for your cipher list in the Cipher List Name field.

Step 6 Click **Add Cipher** to add cipher suites to your cipher list.

Step 7 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 8 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 9 Click **Add** to include the selected cipher suite on your cipher list, or click **Cancel** to leave the list as it is.

Step 10 Repeat Step 6 through Step 9 to add more cipher suites to your list as desired.

Step 11 (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 priorties assigned in this cipher list are only applicable if the cipher list is applied to SSL peering and management services.

Step 12 (Optional) To remove a cipher suite from the list, check the cipher suite's box and then click **Delete**.

Step 13 Click **Submit** when you are done configuring the cipher list.

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group for which you want to manage CA certificates.

Step 3 From the navigation pane, choose **Configure > Security > SSL > Certificate Authorities**.

The SSL CA Certificate List window is displayed (see Figure 12-14).

Figure 12-14 Displaying the SSL CA Certificate List

Step 4 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 5 Add your own CA certificate as follows:

a. Click **Create**. The Creating New CA Certificate window is displayed (see Figure 12-15).

Figure 12-15 Creating a New CA Certificate

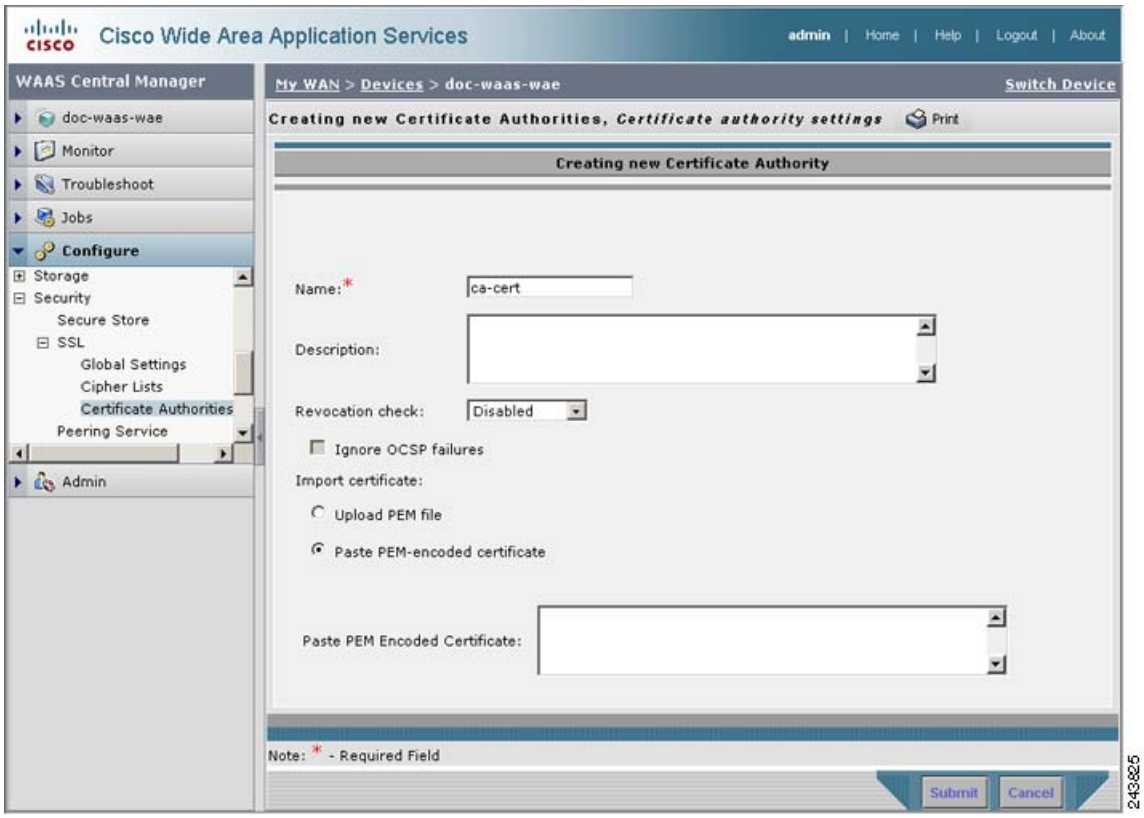

b. Type a name for the certificate in the Certificate Name field.

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 either **Upload PEM File** or **Paste PEM Encoded Certificate**.

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**.

If you are pasting the CA certificate information, paste the text of the PEM format certificate into the Paste PEM Encoded certificate field.

f. Click **Submit** to save your changes.

Step 6 (Optional) To remove a Certificate Authority from the list, select it and then click the **Delete** icon located in the toolbar.

Step 7 Click **Submit** when you are done configuring the CA certificate list.

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 12-5 on page 12-13). 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 GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the device for which you want to configure SSL management services.

Step 3 From the navigation pane, choose **Configure > Security > Management Service**.

The Management Services window is displayed (see Figure 12-16).

Figure 12-16 Configuring SSL Management Services

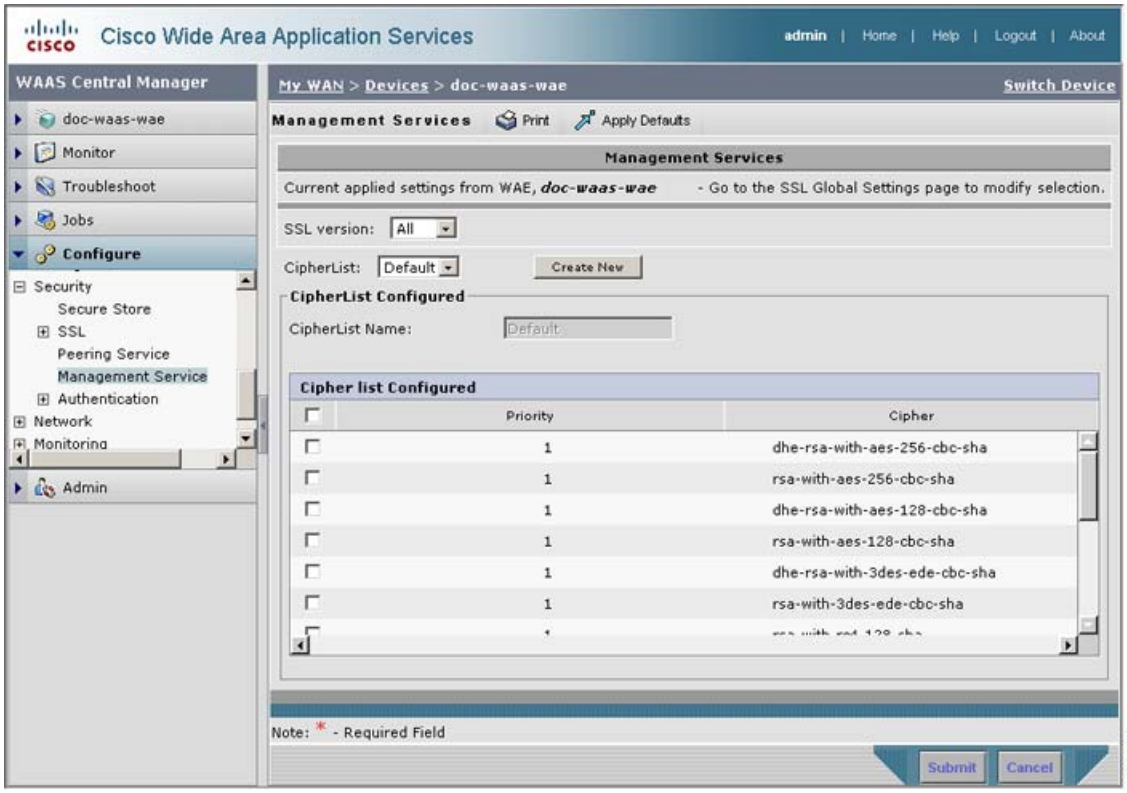

Step 4 In the SSL version field, choose the type of SSL protocol to use. Choose **SSL3** for the SSL version 3 protocol, choose **TLS1** for the Transport Layer Security version 1 protocol, or choose **All** to use both SSL3 and TLS1 SSL protocols.

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 12-2 shows the cipher lists supported with Internet Explorer and Mozilla Firefox:

Table 12-2 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
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 5 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 12-19 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 12-5 on page 12-13). 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 GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the device for which you want to configure SSL peering service.

Step 3 From the navigation pane, choose **Configure > Security > Peering Service**.

The Peering Service window is displayed (see Figure 12-17).

Figure 12-17 Configuring SSL Peering Service

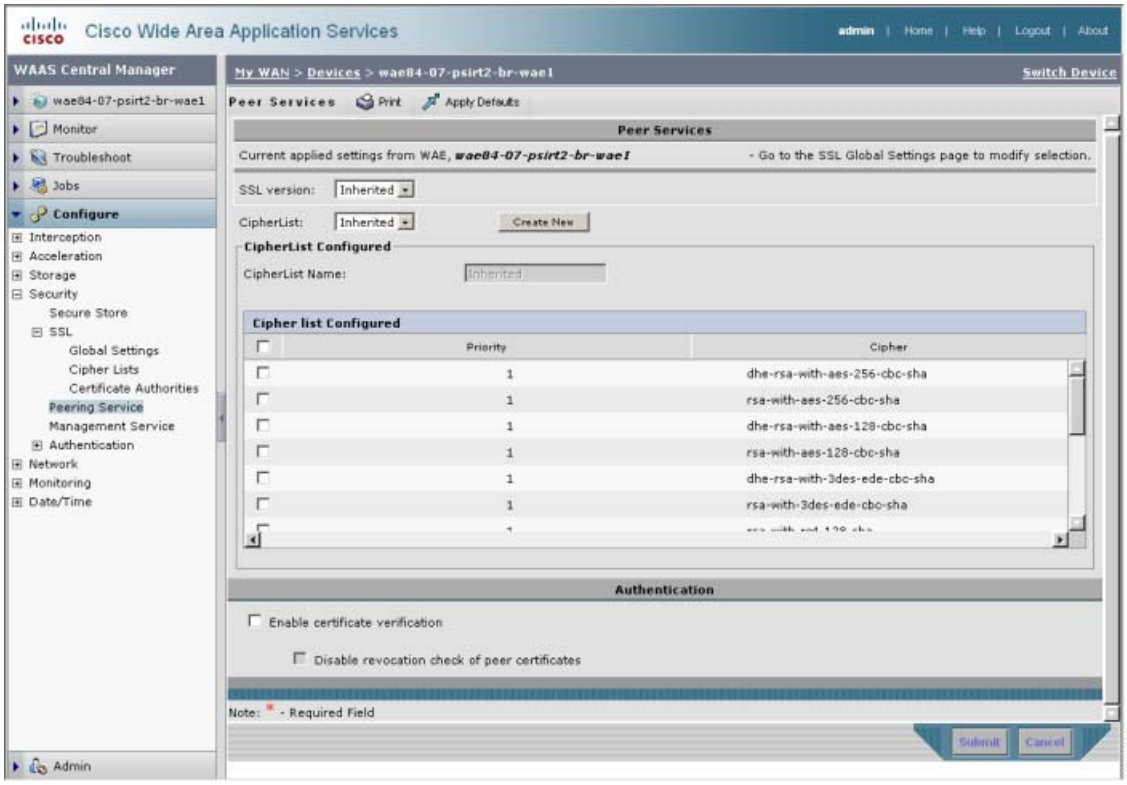

Step 4 In the SSL Version field, choose the type of SSL protocol to use, or choose **Inherited** to use the SSL protocol configured in global SSL settings. Choose **SSL3** for the SSL version 3 protocol, choose **TLS1** for the Transport Layer Security version 1 protocol, or choose **All** to use both SSL3 and TLS1 SSL protocols.

Step 5 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 6 Check the **Disable revocation check for this service** check box to disable OCSP certificate revocation checking.

Step 7 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. See the "Working with Cipher Lists" section on page 12-19 for additional information.

Step 8 Click **Submit**.

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group for which you want to define an accelerated service.

Step 3 From the navigation pane, choose **Configure > Acceleration > SSL Accelerated Services**.

Step 4 To delete an accelerated service, select the service and click **Delete**.

Step 5 Click **Create** to define a new accelerated service. A maximum of 128 accelerated services are allowed. The Basic SSL Accelerated Services Configuration window is displayed (see Figure 12-18).

Figure 12-18 Configuring SSL Accelerated Services—Basic

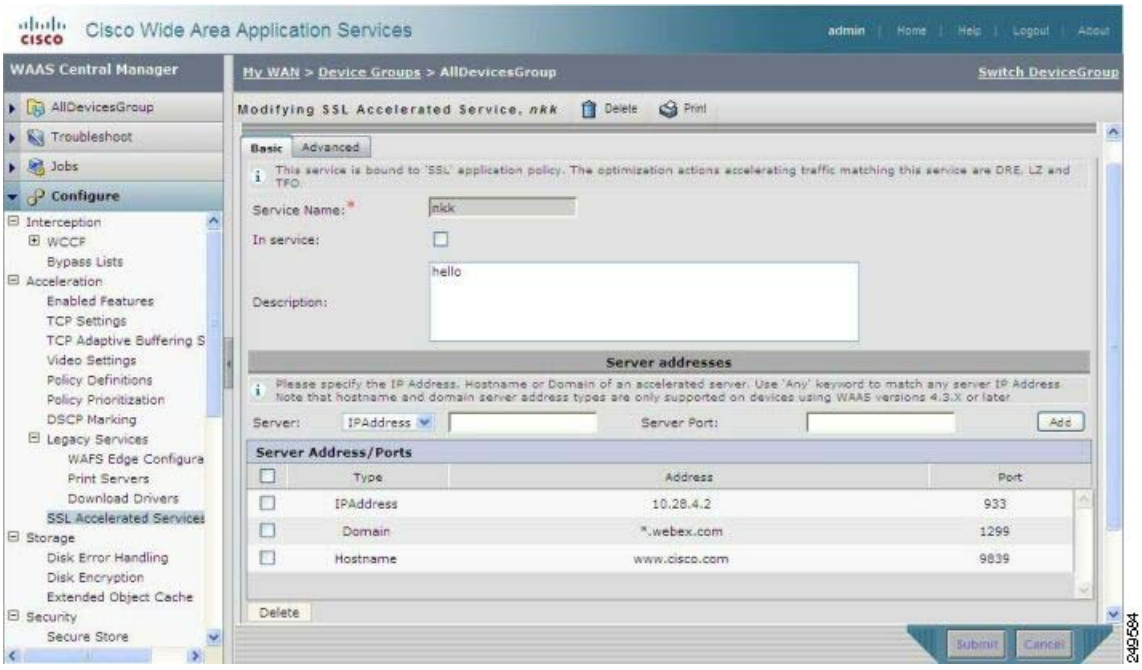

Step 6 Type a name for the service in the Service Name field.

Step 7 To enable this accelerated service, check the **In service** check box.

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 12-19).

Figure 12-19 Configuring Service Certificate and Private Key

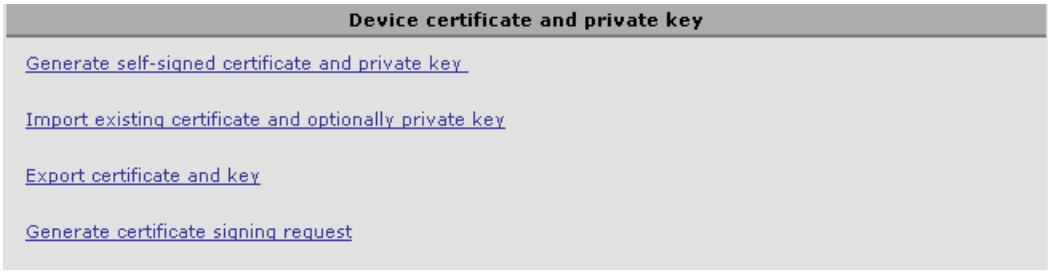

- Click **Generate Self-signed Certificate Key** to have the WAAS devices use a self-signed certificate/key pair for SSL.
- Click **Import Existing Certificate Key** to upload or paste an existing certificate/key pair.
- Click **Export Certificate Key** to export the current certificate/key pair.
- Click **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 12-15.

Step 13 Click the **Advanced Settings** tab to configure SSL parameters for the service. The Advanced SSL Accelerated Services Configuration window is displayed (see Figure 12-20).

Figure 12-20 Configuring SSL Accelerated Services—Advanced

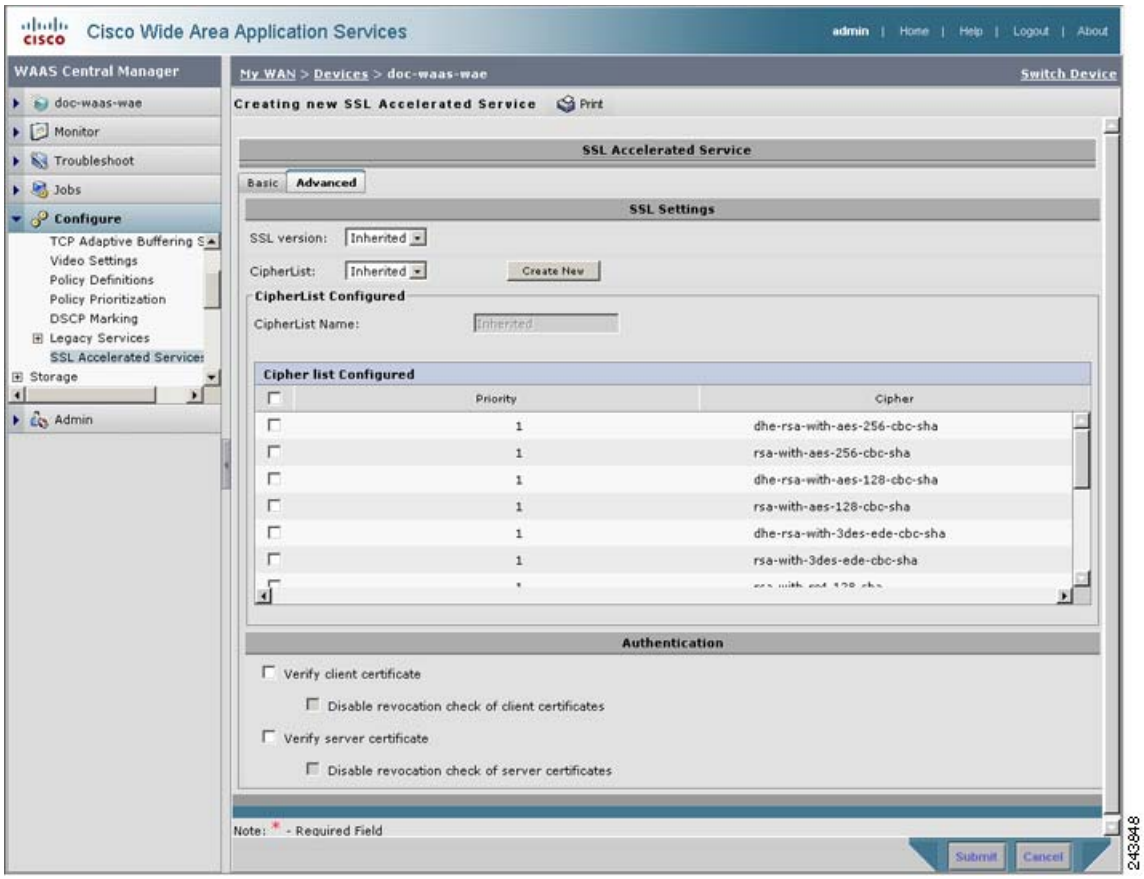

Step 14 (Optional) In the SSL version field, choose the type of SSL protocol to use, or choose **Inherited** to use the SSL protocol configured in global SSL settings. Choose **SSL3** for the SSL version 3 protocol, choose **TLS1** for the Transport Layer Security version 1 protocol, or choose **All** 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 12-19.

Step 16 (Optional) Set the Online Certificate Status Protocol (OCSP) parameters for certificate revocation:

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 Application Policy

Table 12-3 provides an overview of the steps that you must complete to create a new traffic application policy.

Table 12-3 Checklist for Creating a New Application Policy

Task	Additional Information and Instructions
1. Prepare for creating an application policy.	Provides the tasks you need to complete before creating a new application policy on your WAAS devices. For more information, see the "Preparing to Create an Application Policy" section on page 12-29.
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. This step also allows you to assign the application definition to a device or device group. For more information, see the "Creating an Application Definition" section on page 12-30.
3. Create an application 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 classifiers 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 Application Policy" section on page 12-31.

Preparing to Create an Application Policy

Before you create a new application policy, complete the following preparation tasks:

- Review the list of application 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 Application Policies."
- 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 application policy. We recommend you create application policies on device groups so the policy is consistent across multiple WAAS devices.

Creating an Application Definition

The first step in creating an application 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. After creating the application definition, you assign it to a device or device group. 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 GUI navigation pane, choose **Configure > Acceleration > Applications**.

The Applications window is displayed, which displays a list of all applications on your WAAS system. From this window, you can perform the following tasks:

- Click the **Edit** icon next to an application to modify or delete the definition.
- Determine if your WAAS system is collecting statistics on an application. The Monitor Enabled column displays Yes if statistics are being collected for the application.
- Create a new application as described in the steps that follow.

Click the **Create New Application** icon in the taskbar. The Creating Application window is displayed.

Step 2 Enter a name for this application.

The name cannot contain spaces and special characters.

Step 3 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 20 applications, and an error message is displayed if you try to enable statistics for the twenty-first application. 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. However, if you delete an application that you are collecting 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 application policy.

Step 4 (Optional) Enter a comment in the **Comments** field.

The comment you enter appears in the Applications window.

Step 5 Click **Submit**.

The application definition is saved, and options appear in the navigation pane that allow you to assign the application to a device or device group.

Step 6 From the navigation pane, click one of the following options:

- **Assign Device Groups**—Assigns the application to one or more device groups.
- **Assign Devices**—Assigns the application to one or more WAAS devices.

The Device Groups Assignments window or the WAE Assignments window is displayed 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 7 Choose the devices or device groups that you want to assign to this application. To select the devices, use one of the following procedures:

- Click in the taskbar to assign all available WAAS devices or device groups.
- Click next to each WAAS device or device group that you want to assign. The icon changes to when selected. To unassign a device or device group, click the icon again.

Step 8 Click **Submit**.

The icon next to the selected devices changes to , showing that the application has been successfully assigned to the devices.

Creating an Application Policy

After you create an application definition, you need to create an application policy that determines the action a WAAS device takes on the specified traffic. For example, you can create an application 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 application policies on your WAAS system.

The traffic matching rules are contained in the application classifier. These rules, known as match conditions, use Layer 2 and Layer 4 information in the TCP header to identify traffic.

To create an application policy, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group on which you want to create an application policy.

The Device Dashboard window or the Modifying Device Group window is displayed.

Step 3 From the navigation pane, choose **Configure > Acceleration > Policy Definitions**.

The Application Policies window is displayed. (See Figure 12-21.)

Figure 12-21 Application Policies Window

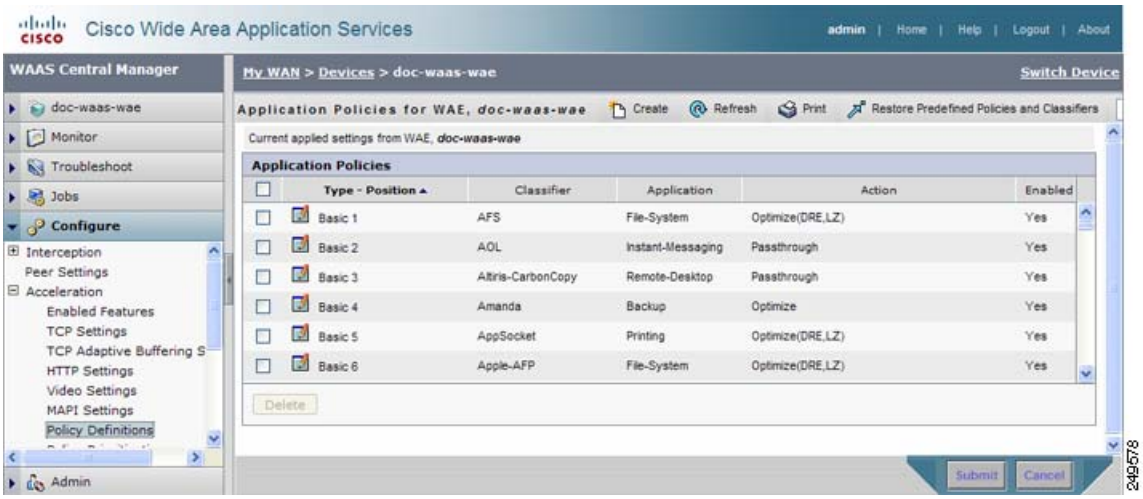

This window displays information about all application policies that reside on the selected device or device group. It shows the type of policy (Basic, WAFS transport, Port Mapper, or Other) and the position of the policy within that type. The position determines the order in which WAAS refers to that policy when determining how to handle application traffic. To change the position of a policy, see the "Modifying the Position of an Application Policy" section on page 12-40. This window also displays the classifier, application definition, and action assigned to each policy.

From the Application Policies window, you can perform the following tasks:

- Place a check next to one or more application policies that you want to delete, and click the **Delete** button to delete the checked policies.
- Click the **Edit** icon next to an application policy to modify or delete that policy.
- Restore predefined policies and classifiers. For more information, see the "Restoring Application Policies and Classifiers" section on page 12-38.
- Create an application policy as described in the steps that follow.

Step 4 Click the **Create New Policy** icon in the taskbar to create a new application policy.

The Creating New Application Policy window is displayed. (See Figure 12-22.)

Figure 12-22 Creating an Application Policy

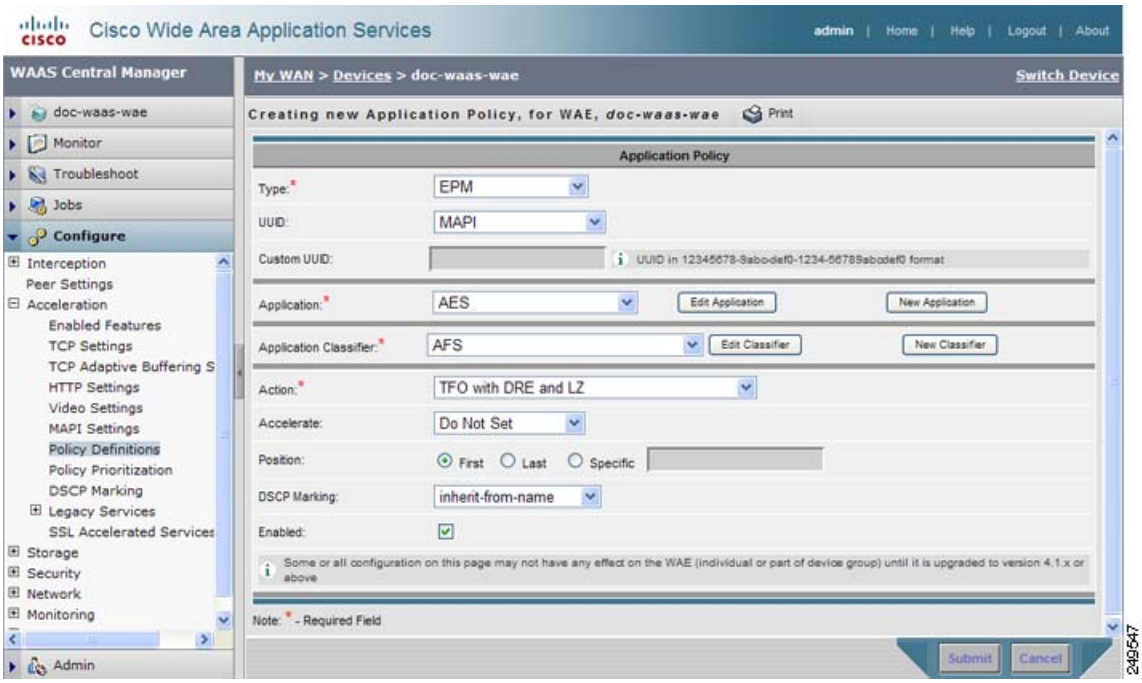

Step 5 From the Type drop-down list, choose the type of application policy.

Table 12-4 describes the types of application policies.

Table 12-4 Application Policy Types

Option	Description
Basic	Standard type of application policy. Choose this option if none of the other types apply.
WAFS Transport	When you enable wide area file services (WAFS), all CIFS traffic going between a branch WAE and a data center WAE is optimized. Choose the **WAFS Transport** option to specify another action (such as passthrough) for CIFS traffic traveling between branch and data center WAEs. For more information on enabling file services, see Chapter 11, "Configuring Wide Area File Services."
EPM	Type of policy for EPM-based applications. EndPoint Mapper (EPM) is a service that dynamically allocates server ports to certain applications. Unlike most applications that always use the same port, applications that rely on the EPM service can be assigned a different port at every request. Because EPM applications do not use a static port, you must specify the application's UUID as a way to identify the application traffic to your WAAS system. When you select the EPM option, the UUID field is enabled so that you can select a preconfigured EPM application or enter the UUID for a custom application.

Step 6 If you chose EPM for the policy type, choose one of the following EPM applications from the **UUID** drop-down list:

- **MAPI**—Uses the predefined UUID associated with the MAPI application, which is a4f1db00-ca47-1067-b31f-00dd010662da.
- **MS-SQL-RPC**—Uses the predefined UUID associated with the SQL Session Manager application, which is 3f99b900-4d87-101b-99b7-aa0004007f07.
- **MS-AD-Replication**—Uses the predefined UUID associated with the Active Directory application, which is e3514235-4b06-11d1-ab04-00c04fc2dcd2.
- **MS-FRS**—Uses the predefined UUID associated with the file replication service, which is f5cc59b4-4264-101a-8c59-08002b2f8426.
- **Custom**—Allows you to enter the UUID for a custom EPM application in the Custom UUID field.

Step 7 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 12-30. This list displays all predefined and new applications on your WAAS system.

  To modify an existing application, choose the application from the drop-down list and click **Edit Application**. You can then change the application's name, add or remove comments, and enable or disable statistics collection for the application. After making the necessary changes, click **Submit** to save your changes and return to the Application Policy window.

- Click **New Application** to create an application. You can specify the application name, enable statistics collection, and specify the DSCP marking value. For the DSCP marking, you can choose to use the global default values (see the "Defining Default DSCP Marking Values" section on page 12-39) or select one of the other defined values. See Table 11-4 on page 11-16 for a description of the supported DSCP marking values. In addition to the values listed in Table 11-4, you can choose copy, which copies the DSCP value from the incoming packet and uses it for the outgoing packet. After specifying the application details, click **Submit** to save the new application and return to the Application Policy window. The new application is automatically assigned to this device or device group.

Step 8 Choose the classifier from the Application Classifier drop-down list to select an existing classifier for this policy.

To modify an existing classifier, select the classifier from the drop-down list and click **Edit Classifier**. You can then change classifier's name, add or remove comments, create a new match condition, or edit the existing match condition. After making the necessary changes, click **Submit** to save your changes and return to the Application Policy window.

Step 9 Click **New Classifier** to create a new classifier for this policy.

The Creating New Application Classifier window then appears so that you can create a new classifier. Complete the following steps to create a new classifier:

a. Enter a name for this application classifier. The name cannot contain spaces or special characters.

b. (Optional) Enter a comment that will appear on the Application Policies window shown in Figure 12-21 on page 12-32.

c. In the Configure Match Conditions section, click the **Create New Match Condition** icon. (If you get a dialog box asking if you want to navigate away from the page, click **OK**.) The Creating New Match Condition window is displayed. (See Figure 12-23.)

Figure 12-23 Creating a New Match Condition

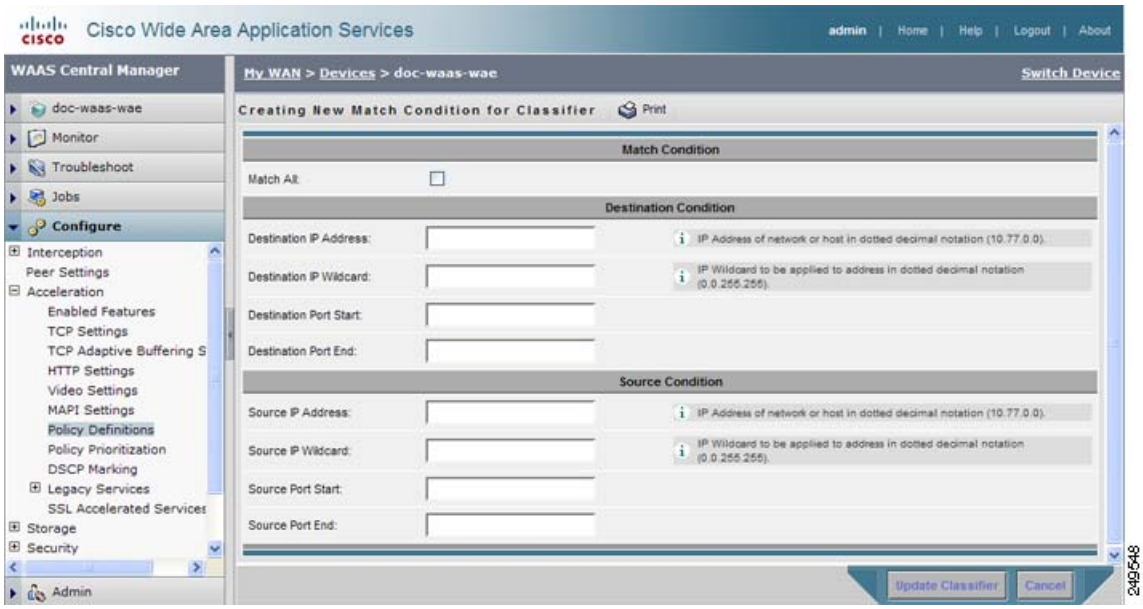

d. Check the **Match All** check box to create a condition that matches all traffic. Checking the Match All check box automatically disables all other fields in the window.

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 **Update Classifier**. You return to the Creating New Application Classifier window. The new match condition appears at the bottom of this window.

g. Click **Submit**. You return to the Creating New Application Policy window.

Step 10 From the Action drop-down list, choose the action that your WAAS device should take on the defined traffic. Table 12-5 describes each action.

Table 12-5 Action Descriptions

Action	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 Data Redundancy Elimination	Applies both TFO and data redundancy elimination (DRE) to matching traffic. DRE removes redundant information before sending the shortened data stream over the WAN. DRE operates on large data streams (tens to hundreds of bytes or more).
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 and LZ	Applies TFO, DRE, and LZ compression to matching traffic.

Step 11 From the Accelerate drop-down list, choose one of the following additional acceleration actions that your WAAS device should take on the defined traffic:

- **Do Not Set**—No additional acceleration is done.
- **MS Port Mapper**—Accelerate using the Microsoft Endpoint Port Mapper (EPM).
- **CIFS**—Accelerate using the CIFS Accelerator.
- **HTTP**—Accelerate using the HTTP Accelerator.
- **NFS** —Accelerate using the NFS Accelerator.
- **MAPI**—Accelerate using the MAPI Accelerator.
- **VIDEO** —Accelerate using the Video Accelerator.

Step 12 Choose one of the following positions for this application policy by click the appropriate Position radio button:

- **First**—Places this policy at the top of the position list so that the WAAS device tries to classify traffic using this policy before moving onto the second policy in the list. If you already have a policy in the first position, that policy moves down to number two in the list.
- **Last**—Places this policy at the bottom of the position list, making it the last policy that the WAAS device uses to classify traffic. If you already have a policy in the last position, that policy becomes the second to last in the list.

  If a device goes through all the policies in the list without making a match, then the WAAS device passes through the traffic unoptimized.
- **Specific**—Allows you to enter a specific position for this policy. If you already have a policy in the specified position, that policy moves down one in the list.

Step 13 (Optional) Choose a value from the DSCP Marking drop-down list and see Table 11-4 on page 11-16 for a description of the supported values. In addition to the values listed in Table 11-4, 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.

The DSCP value set at the application level applies to all classifiers associated with the application. If you set a DSCP value in a policy, it overrides the DSCP value set at the application or global level.

Step 14 Check the **Enabled** check box to activate this policy. To disable this policy, uncheck this box.

Step 15 Click **Submit**.

The new policy appears in the Application Policies window. (See Figure 12-21 on page 12-32.)

Managing Application Acceleration

This section contains the following topics:

- Viewing a List of Applications, page 12-37
- Viewing a Policy Report, page 12-38
- Viewing a Classifier Report, page 12-38
- Restoring Application Policies and Classifiers, page 12-38
- Monitoring Applications, page 12-39
- Defining Default DSCP Marking Values, page 12-39
- Modifying the Position of an Application Policy, page 12-40
- Modifying the Acceleration TCP Settings, page 12-41

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group on which you want to view applications.

Step 3 From the navigation pane, choose **Configure > Acceleration > Policy Definitions**. The Application Policies window is displayed.

Step 4 Click the Application column header to sort the column by application name so you can more easily locate a specific application.

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 GUI navigation pane, choose **Configure > Acceleration > Policies**.

The policy report appears. It lists each device or device group and the number of active policies on the device or device group.

Step 2 Click the **Edit** icon next to a device or group to see the application policies that are defined on it.

Viewing a Classifier Report

To view a report of the classifiers that reside on each WAE device or device group, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Configure > Acceleration > Classifiers**.

The classifier report appears. It lists each classifier that is defined, and the number of devices on which it is configured.

Step 2 Click the **View** icon next to a classifier to see a report of the devices and device groups on which the classifier is configured.

Step 3 Click the **Edit** icon next to a device or group to see the application policies that are defined on it.

Restoring Application Policies and Classifiers

The WAAS system allows you to restore the predefined policies and classifiers that shipped with the WAAS system. For a list of the predefined policies, see Appendix A, "Predefined Application Policies."

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 classifiers, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or group on which you want to restore policies.

Step 3 From the navigation pane, choose **Configure > Acceleration > Policy Definitions**.

The Application Policies window is displayed.

Step 4 Click the **Restore Predefined Policies and Classifiers** taskbar icon to restore over 150 policies and classifiers 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

After you create an application 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 12-30.

You can use the Traffic Optimization report to monitor a specific application. For more information, see the "Optimization Summary Report" section on page 16-39.

Defining Default DSCP Marking Values

According to policies that you define in an application definition and an application 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. This value applies to the traffic if a lower level value is not defined.
- Application—You can define the DSCP value in an application definition at the device or device group level but not at the global application definition level. This value applies to all traffic associated with the application on a particular device or device group and overrides the global default.
- Policy—You can define the DSCP value in an application policy. This value applies only to traffic that matches the classifiers 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 12-39

Defining the Default DSCP Marking Value

To define the global default DSCP marking value, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or group where you want to define the default DSCP marking value.

Step 3 From the navigation pane, choose **Configure > Acceleration > DSCP Marking**. The Global DSCP Settings window is displayed.

Step 4 Choose a value from the Global Default DSCP Marking drop-down list and see Table 11-4 on page 11-16 for a description of the supported values. In addition to the values listed in Table 11-4, the default setting is copy, which copies the DSCP value from the incoming packet and uses it for the outgoing packet.

Step 5 Click **Submit** to save the settings.

Modifying the Position of an Application Policy

Each application 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.

For information on how to assign a position to a new policy, see the "Creating an Application Policy" section on page 12-31.

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 application 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 16, "Monitoring and Troubleshooting Your WAAS Network."

To modify the position of an application policy, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or group that contains the application policy to modify.

Step 3 From the navigation pane, choose **Configure > Acceleration > Policy Prioritization**.

Step 4 The Application Policies window is displayed. This window categorizes policies into these categories: Basic, Other, Port Mapper, and WAFS.

Step 5 Click the arrow next to the appropriate category to display the list of applications for that category. (See Figure 12-24.)

In most cases, the application you want to change the position for will be located under the Basic Policies category because that category contains a majority of the predefined applications that shipped with the WAAS system. For a list of these predefined policies, see Appendix A, "Predefined Application Policies."

Figure 12-24 Modifying the Position of Application Policies

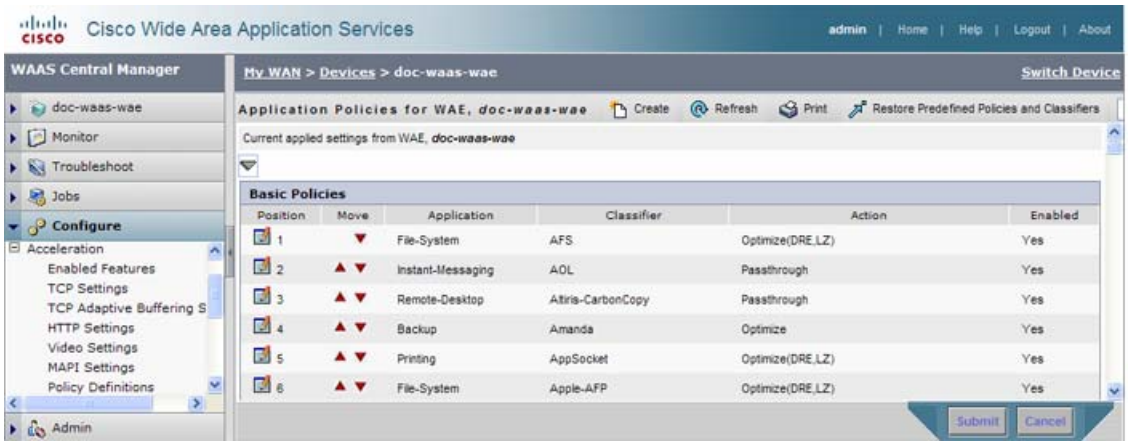

Step 6 Click the arrow next to the policy category to view the list of applications for that category.

Step 7 Use the up and down arrows (▲ ▼) next to a policy to move that policy higher or lower in the list.

Step 8 If you determine that a policy is not needed, follow these steps to delete the policy:

a. Click the **Edit** icon next to the policy you want to delete.

The Modifying Application Policy window is displayed.

b. Click the **Delete** icon in the taskbar.

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 12-43.

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 12-43.

To modify the acceleration TCP settings, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group for which you want to change the acceleration TCP settings.

Step 3 From the navigation pane, choose **Configure > Acceleration > TCP Settings**. The Acceleration TCP Settings window is displayed.

Step 4 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 5 Modify the TCP acceleration settings as needed. See Table 12-6 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 12-43.

Table 12-6 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 6 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 12-43.

Step 7 Click **Submit**.

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-511/512—Default BDP is 32 KB
- WAE-611/612—Default BDP is 512 KB
- WAE-674 —Default BDP is 2048 KB
- WAE-7326 —Default BDP is 2048 KB
- WAE-7341 —Default BDP is 2048 KB
- WAE-7371 —Default BDP is 2048 KB
- WAVE-274 —Default BDP is 2048 KB
- WAVE-474 —Default BDP is 2048 KB
- WAVE-574—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 12-43.

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or device group) for which you want to configure the TCP adaptive buffering settings.

Step 3 In the navigation pane, choose **Configure > Acceleration > TCP Adaptive Buffering Settings**. The TCP Adaptive Buffering Settings window is displayed.

Step 4 To enable TCP adaptive buffering, check the **Enable** check box. The default is enabled.

Step 5 In the Send Buffer Size and Receive Buffer Size fields, enter the maximum size in kilobytes of the send and receive buffers.

Step 6 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 13

Configuring and Managing WAAS Legacy Print Services

This chapter describes how to configure and manage the WAAS legacy print services feature that allows WAEs to function as print servers in your branch offices.

For information on other print services options, see the "WAAS Print Services" section on page 1-8.

Note The legacy print services feature is deprecated in WAAS version 4.2.1. It still operates but is not supported and will be removed in a future version. If you enable legacy print services, an alarm is raised on the device and you are alerted in all Central Manager GUI pages and the CLI if you configure any legacy print services settings. Legacy print services users should migrate to the Windows Print accelerator.

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 WAAS Print Services, page 13-1
- Planning for Print Services, page 13-5
- Configuring Print Services, page 13-7
- Managing Print Services, page 13-25
- Troubleshooting Print Services, page 13-34

About WAAS Print Services

The WAAS software includes print services that allow you to change a branch WAE into a print server. This functionality eliminates the need for a separate print server in the branch office. The WAAS print services feature is available for Windows clients and works with any IP-based network printer.

The WAAS Central Manager GUI allows you to distribute print drivers to specific WAAS print servers on a per-device or a per-device group basis. You can also manage the print queue and monitor the status of a print job from the Print Services Administration GUI that you can open from the WAAS Central Manager GUI.

This section contains the following topics:

- Branch Office Printing Topology, page 13-2
- WAAS Print Services, page 13-3

Branch Office Printing Topology

A typical print server topology in a branch office consists of several client desktops using a single print server to proxy for multiple printers. The WAAS legacy print services solution allows you to host the print server on the branch WAE along with the WAAS software. Figure 13-1 shows the WAAS print services topology for branch offices and data centers.

Figure 13-1 Branch Office WAAS Print Services Topology

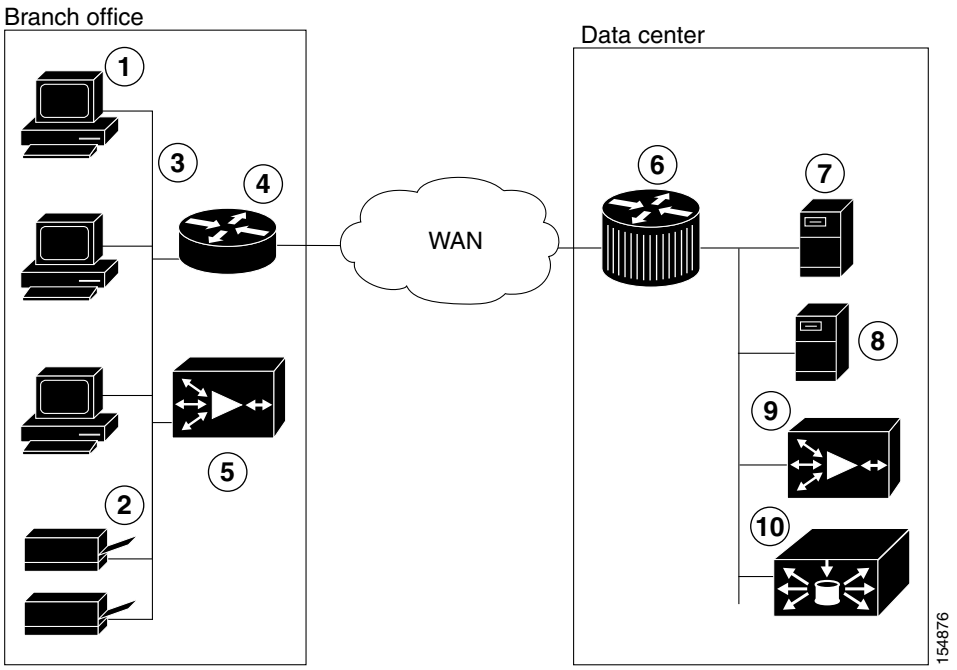

1	CIFS clients in the branch office	**6**	Cisco router in the data center
2	Networked printers in the branch office	**7**	File server in the data center
3	Local LAN in the branch office	**8**	Backup file server in the data center
4	Cisco router in the branch office	**9**	Data Center WAE
5	WAE (WAAS print server) in the branch office	**10**	WAAS Central Manager in the data center

In this configuration, the branch WAE provides local print services to branch clients. Microsoft clients use the WAAS print server in the same way that they may currently be using a Microsoft print server:

- Clients can add or delete print queues to their local computers by using the Windows Wizard.
- Clients can spool print jobs to the WAAS print server.

- The WAAS print server communicates with the printer for print functions.

WAAS Print Services

The WAAS print services solution on the branch WAE consists of two major components:

- Samba—WAAS uses Samba to enable Microsoft clients to add and delete print queues, add and delete drivers, browse for print queues, and spool jobs to the WAAS print server.
- Common Unix Printing System (CUPS)—WAAS uses CUPS for print queue management and for sending the Microsoft clients' spooled print jobs to the appropriate printers using TCP/IP.

This section provides information about the WAAS print services and contains the following topics:

- Print Driver Support and Interoperability, page 13-3
- Printer Clustering, page 13-4
- Print Services Users, page 13-4
- Feature Support, page 13-4

Print Driver Support and Interoperability

WAAS WAE incorporates a Print Server based on the integration of open source Samba and CUPS technology. During the testing process, it has been determined that certain Print Drivers with complex features, such as sophisticated paper handling, may not be Point-and-Print compatible with WAAS. Most notably, Fiery Drivers incorporated into some Printer Manufacturer solutions are not compatible with Samba. Other Multi Function Printers (MFP) may also have limited functionality when working with Samba and are not supported by WAAS.

To determine if a Print Driver is compatible with WAAS, perform the Add Driver processes with a WAE using the Add Printer Wizard. Compare all the client Print features available after creating a print queue and compare it to a similar installation on a Microsoft Windows Print Server. If there are obvious feature inconsistencies, it is indicative of a Print Driver that cannot be used with WAAS Print Server for Point-and-Print. As a workaround, an installation on each client desktop from a CD or other source will be required.

When using the WAAS print services in a Windows XP Pro/Windows 2003 Server environment, you must register the WAE with Active Directory for the automatic printer driver download feature to operate correctly. This is due to a default computer policy for domain members that does not allow the host to download drivers from an unregistered device. A user will see a message similar to the following when encountering this issue: "A policy is in effect on your computer which prevents you from connecting to this print queue. Please contact your system administrator."

Additionally, note the following:

- Printing is possible only when the branch WAE is running.
- The WAAS print solution does not offer authentication. Any user may access and send print jobs to the WAAS print server.
- Only IP-based network printers are supported in WAAS. We recommend that printers directly attached to the WAAS print server use parallel, serial, or USB IP adapters, such as those available from HP or Linksys.
- WAAS supports Raw Queue so that all rendering of files occurs on the client machine using vendor-provided printer drivers, and the print server performs Win2003-style print queue management.

- WAAS supports 32-bit drivers. The Samba version used in the WAAS software does not support 64-bit print drivers.

Printer Clustering

Printer clustering enables administrators to group several printers together to provide failover and load-balancing capabilities. You can include as many printers as you want in a printer cluster, but we recommend that you include no more than 12 printers. Only printers of the same model and capabilities should be grouped together. Printers in a cluster should be geographically close to each other. The default spooling space for all printers (including the printer cluster) is 1 GB.

The printer cluster appears as a single print queue to Microsoft clients using the WAAS print server. The WAAS print server does not support the selection of a default printer within a cluster. When you send a print job to a printer in a printer cluster, the print job is automatically forwarded by the WAAS print server, which then sends it to the first available printer.

Print Services Users

The following types of users exist in the WAAS print services environment:

- Administrative User—A user who can add and modify printer information and membership in printer clusters.
- Print User—A user who can print a job from a printer, but who cannot make any additions or changes to the printer configuration. A print user can delete and pause their own print jobs.

Note Job owners can manage their own job in the Microsoft client and Print Services Administration GUI while the job is still spooling.

Feature Support

Various print services features are handled by the various software components, such as Samba, CUPS, and the WAAS software. Table 13-1 shows which tool provides the capability for a specific print services feature.

Table 13-1 WAAS Print Services Feature Support

Function	Software that Provides the Function
Add, modify, or delete printers	CUPS
Add, modify, or delete printer cluster (called classes in CUPS)	CUPS
View and control print jobs	CUPS, Windows printer queue console
Monitor status of individual printers	CUPS
View list of printers attached to a WAAS print server, the drives associated with a specific printer, and the list of all distributed drivers	WAAS Central Manager GUI

Table 13-1 WAAS Print Services Feature Support (continued)

Function	Software that Provides the Function
Diagnostics and troubleshooting	CUPS and Samba perform their diagnostics. The WAAS Central Manager GUI displays error messages relating to driver distribution.
Client printer driver installation from print server	Samba[1]
Print driver distribution to print servers	WAAS Central Manager GUI
Advertising printer availability through Samba and Windows Active Directory	Samba
Log export	Manual FTP using the CLI
Admin authentication	WAAS Central Manager
Secure administration	WAAS software

1. The Samba version used in the WAAS software does not support 64-bit print drivers.

Planning for Print Services

This section describes the information that you should collect before you begin configuring print services. You need the information outlined in this section whether you are setting up WAAS print services for the first time, or migrating from Microsoft print services. When planning for print services, these steps will help avoid configuration errors:

1. Plan for your printers' network capabilities and driver support.

   This information includes the printer type (Postscript, PCL, or other), the printer network protocol, and port (LPD, IPP, LPQ, LPR and port), and whether printers should be clustered together for failover. To obtain driver information, find either CDs or servers containing the printer drivers to be installed on the WAAS print server for support of point-and-print functionality.

2. Plan for print queue configuration.

   This information includes the names and types of print queues to be created for a print server, and how they interact with the printers' networking capabilities and ports.

The following sections provide more information on planning for print services:

- Identifying the Print Administration Users, page 13-5
- Obtaining Printer Information, page 13-6
- Planning Worksheets, page 13-6

Identifying the Print Administration Users

All your branch office clients can use the print services on a WAAS print server, but only administrative users can perform printer management tasks such as managing print queues. These administrative users need to have admin accounts created on the branch WAE and on the WAAS Central Manager device.

Identify the users that you want to have print administrator privileges for because you will need to create these users on the WAAS print server and on the WAAS Central Manager device.

Note We recommend limiting the number of admin users who have access to your printers. Each admin user who accesses a printer's properties causes the printer object to be updated, which forces all clients to update their copy of the printer object on their next session. If admin users access printers on a regular basis, delays and increased traffic can result.

Obtaining Printer Information

Before configuring print services, it is useful to obtain the following information about each of the printers on your network:

- Printer type
- Printer protocol
- Number and name of print queues
- Printer connectivity (direct network or using a print server gateway device)
- Printer drivers

  Drivers may reside on a vendor-provided CD-ROM. Commonly used drivers are also bundled with the Windows operating system. To access these common drivers, open the Add Printer Wizard on a Windows system.

Planning Worksheets

You can use Table 13-2, Table 13-3, and Table 13-4 as planning worksheets to help ensure a successful configuration of the WAAS print services.

Table 13-2 Security Model and Directory Services

Description	Value
Print server IP	
WINS server IP	
Print server NetBIOS name (optional)[1]	
Print server admin workstation user name	

1. If the NetBIOS name is not specified for the print server, the hostname is used instead. If the hostname is longer than fifteen characters, the hostname will be truncated to fifteen characters.

Table 13-3 Network Capabilities and Driver Support

Description	Value
Printer type (PostScript, PCL, other)	
Printer network protocol (LPD, IPP, LPQ, LPR)	
Printer port	
Printer cluster	
Driver file location and name	

Table 13-4 Print Queue Configuration

Description	Value
Print queue name	
Print queue type	

Configuring Print Services

This section describes how to configure print services in your WAAS network.

Configuration Checklist

Note The legacy print services feature is deprecated in WAAS version 4.2.1. It still operates but is not supported and will be removed in a future version. If you enable legacy print services, an alarm is raised on the device and you are alerted in all Central Manager GUI pages and the CLI if you configure any legacy print services settings. Legacy print services users should migrate to the Windows Print accelerator.

Table 13-5 lists the configuration steps you must complete to set up print services in your WAAS network.

Table 13-5 Checklist for Configuring WAAS Print Services

Task	Additional Information and Instructions
1. Review the information in the planning section.	Before configuring print services in your WAAS network, see the "Planning for Print Services" section on page 13-5.
2. Prepare your WAE device and WAAS Central Manager device for print services.	To prepare your WAAS system for print services, you need to specify the WINS server name and NetBIOS name for your WAE device and WAAS Central Manager device. In addition, you also need to enable either the CIFS accelerator or the legacy WAFS edge services on the WAE device. For more information, see the "Preparing your WAE Device and Central Manager for Print Services" section on page 13-8.
3. Create accounts with print admin privileges on the WAE device and the WAAS Central Manager device.	To configure print services using the Print Services Administration GUI and to install drivers in the central repository on the WAAS Central Manager device, you need to create an accounts with print admin privileges on each of these devices. For more information, see the "Creating Accounts with Print Admin Privileges" section on page 13-10.
4. Enable print services.	By default, the WAAS print services feature is disabled on all WAAS devices. To enable print services, see the "Enabling Print Services" section on page 13-11.
5. Add a printer to the WAAS print server.	To add printers to your new WAAS print server, see the "Adding a Printer to the WAAS Print Server" section on page 13-12.
6. Create print clusters (optional).	To provide failover functionality to print services, you can combine multiple printers into a print cluster. For more information, see the "Adding Printer Clusters" section on page 13-14.

Table 13-5 Checklist for Configuring WAAS Print Services (continued)

Task	Additional Information and Instructions
7. Set up the WAAS Central Manager to be the central driver repository.	WAAS allows you to centrally manage all your print drivers from the WAAS Central Manager if it is set up as the main driver repository. For more information, see the "Setting Up the WAAS Central Manager as the Driver Repository" section on page 13-16.
8. Distribute drivers to the WAAS print servers.	After you have set up the WAAS Central Manager as the driver repository and have loaded drives into the repository, you can distribute the drivers to all the print servers in your WAAS network. For more information, see the "Distributing Drivers to the WAAS Print Servers" section on page 13-20.
9. Associate the print driver with your printer.	After you distribute a driver to a WAAS print server, you need to associate the driver with the correct printer on the WAAS print server. For more information, see the "Associating a Driver with your Printer" section on page 13-22.
10. Initialize each print driver on your WAAS print servers.	To prevent printing issues, you need to initialize each print driver before they are installed and used by your branch office clients. For more information, see the "Initializing Print Drivers" section on page 13-23.
11. Add printers to your branch office clients.	After the print services have been properly configured, your branch office clients need to add printers using the Microsoft Add Printers wizard. For more information, see the "Adding the WAAS Print Server to Your Branch Office Clients" section on page 13-24.

Preparing your WAE Device and Central Manager for Print Services

Before you enable print services, you need to prepare your WAE device and Central Manager by specifying the WINS server name and the NetBIOS name. In addition, you also need to enable either the CIFS accelerator or the legacy WAFS edge services on the WAE device. In some cases, these settings may already have been configured when the WAE device and Central Manager were first set up.

Note If you are not using WINS, you can still deploy legacy print services by manually adding the printer to Active Directory. This makes it possible for clients to find the printer, since it will not appear in the network browser, unless there is a Windows client deployed in the same subnet as the print server.

To prepare your WAE device and Central Manager for print services, follow these steps:

Step 1 (Optional) Specify the NetBIOS name of the WAE device:

a. From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

b. Click the **Edit** icon next to the WAE device on which you want to enable print services.

c. From the navigation pane, choose *Device Name* **> Activation**. The Device Activation window appears.

d. Enter the NetBIOS name of the WAE device in the NetBIOS Name field.

Note If the WAE is operating in nontransparent mode, you must configure identical names for the NetBIOS name and the hostname of the device that you enter in the Name field.

e. Click **Submit** to save your changes.

Once you enable print services on this WAE device, the NetBIOS name is used as the Print Server name for the device. If you do not specify a NetBIOS name, the hostname of the device is used instead. The hostname of the device is specified on the Activation page, described in the "Modifying Device Properties" section on page 9-1.

To use the CLI to specify the NetBIOS name, use the **windows-domain netbios-name** global configuration command.

Step 2 Specify the NetBIOS name of WAAS Central Manager device:

a. From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

b. Click the **Edit** icon next to the WAAS Central Manager device.

c. From the navigation pane, choose *Device Name* **> Activation**. The Device Activation window appears.

d. Enter the NetBIOS name of the WAAS Central Manager in the NetBIOS Name field.

e. Click **Submit** to save your changes.

Step 3 Specify the WINS server name of the WAE device:

a. From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

b. Click the **Edit** icon next to the WAE device on which you want to enable print services.

c. From the navigation pane, choose **Configure > Network > WINS**.

d. Enter the name of the WINS server in the WINS Server field.

e. Click **Submit** to save your changes.

To use the CLI to specify the NetBIOS name, use the **windows-domain wins-server** global configuration command.

If you are not using WINS, manually add the printer to Active Directory.

Step 4 Specify the WINS server name of the WAAS Central Manager device:

a. From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

b. Click the **Edit** icon next to the WAAS Central Manager device.

c. From the navigation pane, choose **Configure > Network > WINS**.

d. Enter the name of the WINS server in the WINS Server field.

e. Click **Submit** to save your changes.

Step 5 Enable the CIFS accelerator on the WAE device (or skip to Step 6):

a. From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

b. Click the **Edit** icon next to the WAE device on which you want to enable print services.

c. From the navigation pane, choose **Configure > Acceleration > Enabled Features**.

d. Check the **CIFS Accelerator** check box.

e. Click **Submit** to save your changes.

Step 6 Enable legacy WAFS edge services on the WAE device (or perform Step 5):

a. From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

b. Click the **Edit** icon next to the WAE device on which you want to enable print services.

c. From the navigation pane, choose **Configure > Acceleration > Legacy Services > Edge Configuration**.

d. Check the **Enable Edge Server** check box. The CIFS and QoS configuration settings become enabled. For information on configuring these settings, see the "Configuring the Edge Devices" section on page 11-14.

e. Click **Submit** to save your changes.

Note You must reboot or reload the device after enabling edge services.

Creating Accounts with Print Admin Privileges

To configure print services using the Print Services Administration GUI, you need to create an account with print admin privileges on the WAAS Central Manager device. This account also lets you upload drivers to the central repository.

You also must assign a domain to any user account that has print admin privileges. A domain defines the device groups or WAEs that are accessible by the user. All WAEs to which the user needs access should be members in the assigned domain.

To create an account with print admin privileges on the WAAS Central Manager device, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, 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.

Step 3 In the Username field, enter the user account name. User names are case-sensitive and support special characters.

Step 4 Check the **Local User** check box.

Step 5 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.

Step 6 From the CLI Privilege Level drop-down list, choose either **0 (normal user)** or **15 (super user)**.

Step 7 Check the **Print Admin** check box.

Step 8 (Optional) Fill in the fields in the User Information and Comments section.

Step 9 Click **Submit** to save your changes.

Next you must assign the print role to the account.

Step 10 Click the **Role Management** tab.

The Role Management for User Account window appears with all configured role names listed.

Step 11 Click the **Assign** icon (blue cross mark) that appears next to the print role.

Step 12 Click **Submit**.

Next, you must assign a domain to the account.

Step 13 Click the **Domain Management** tab.

The Domain Management for User Account window appears with all configured domain names listed.

Step 14 Click the **Assign** icon (blue cross mark) that appears next to the domain that you want to assign to the account. All WAEs to which the user needs access should be members in the assigned domain.

Step 15 Click **Submit**.

For more information about creating accounts on the WAAS Central Manager device, see the "Creating a New Account" section on page 7-4.

To create an account with print admin privileges (privilege level 15) from the CLI, you can use the **username** global configuration command.

Enabling Print Services

You must use the WAE Device Manager GUI to enable print services on a WAE.

Note The legacy print services feature is deprecated in WAAS version 4.2.1. It still operates but is not supported and will be removed in a future version. If you enable legacy print services, an alarm is raised on the device and you are alerted in all Central Manager GUI pages and the CLI if you configure any legacy print services settings. Legacy print services users should migrate to the Windows Print accelerator.

You can enable print services on any WAE, but you cannot enable print services on the WAAS Central Manager. Once you enable print services, the branch WAE becomes a WAAS print server and an anonymous FTP user account is created on the device. This FTP account is used to access print drivers on the WAAS Central Manager.

To enable print services on a branch WAE, follow these steps:

Step 1 Log into the WAE Device Manager GUI by going to the following URL:

https://*wae-ip-address*:8443/mgr

wae-ip-address is the IP address of the branch WAE that you want to enable with print services.

Step 2 Choose **Configuration** from the Cisco WAE menu, and click the **Print Services** tab.

The Print Services window appears. (See Figure 13-2.)

Figure 13-2 Print Services Window

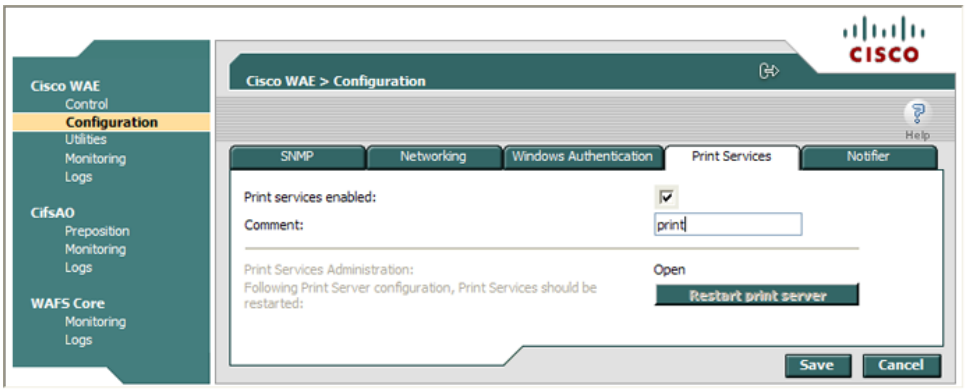

Step 3 Check the **Print services enabled** check box to enable print services.

Step 4 (Optional) Enter a comment in the Comment field.

This comment string is visible next to the print server when users are browsing in Windows Explorer.

Step 5 Click **Save** to restart print services.

Upon restart, the Open link for the Print Services Administration GUI is enabled. Once print services is enabled on the device, the device becomes a WAAS print server.

Note The **Restart Print Server** button is used when a change has been performed through the CLI and the print server must be restarted.

To enable print services from the CLI, you can use the **print-services enable** global configuration command.

Adding a Printer to the WAAS Print Server

After you enable print services, you must add one or more printers to the new WAAS print server. When you add a printer, a print queue is also added that is used by Windows print clients to spool their jobs to the WAAS print server.

To add a printer to the WAAS print server, follow these steps:

Step 1 Open the Print Services Administration GUI by doing either of the following:

- From the WAAS Central Manager GUI navigation pane, choose **Configure > Legacy Services > Print Servers**, click the **Edit** icon next to the name of the print server that you want to add to a cluster, and then click the **CUPS** icon in the taskbar.
- From the WAE Device Manager GUI, choose **Configuration** from the Cisco WAE menu, click the **Print Services** tab, and then click the **Open** link.

The Print Services Administration GUI opens. (See Figure 13-3.)

Figure 13-3 Print Services Administration GUI

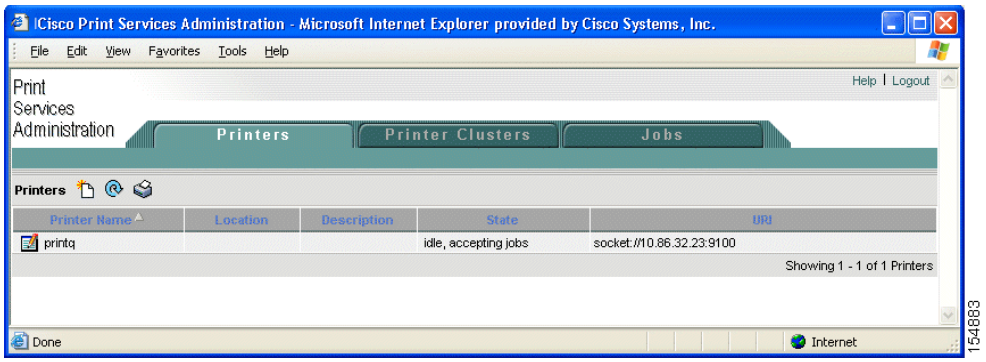

Step 2 Click the **Add Printer** icon. You are prompted to enter a username and password.

Step 3 Enter the username and password for the print admin account you created on this WAE device, and click **OK**.

The Add Printer window appears. (See Figure 13-4.) For more information on creating this account, see the "Creating a New Account" section on page 7-4.

Figure 13-4 Adding a Printer to the WAAS Print Server

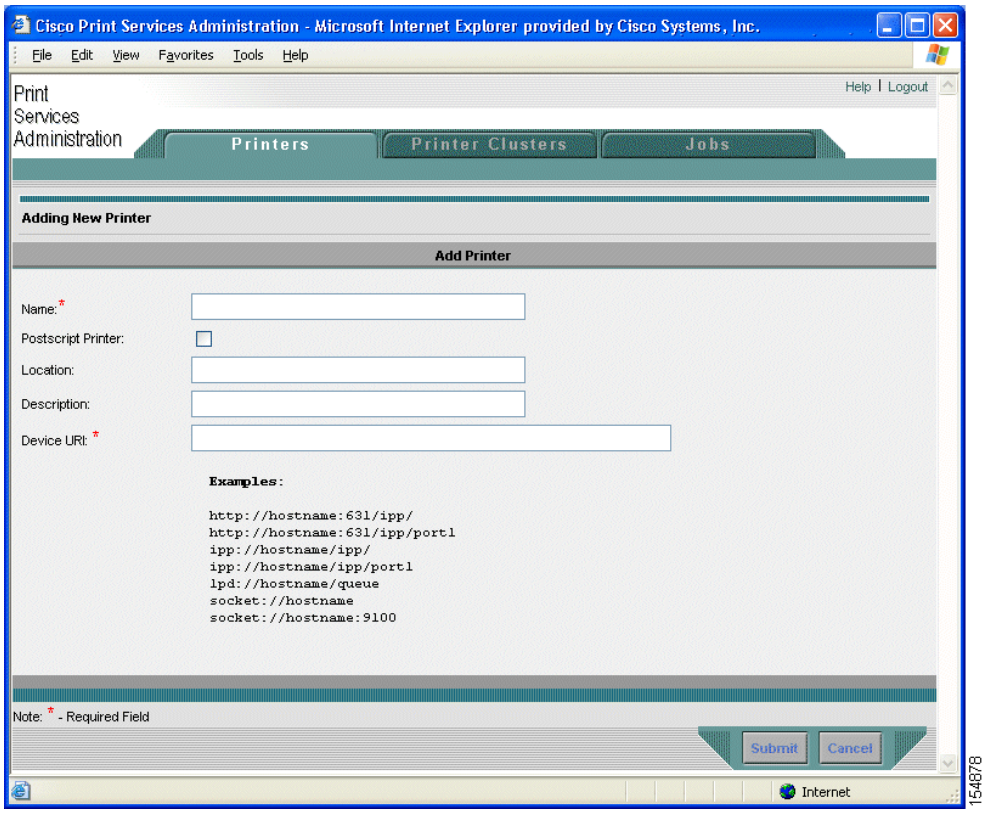

Step 4 Enter the following information in the provided fields:

- **Printer name** (required)—User-defined name of the printer (maximum of 127 characters) that can contain only letters, numbers, and underscore. No spaces are allowed and the name must be unique across the system. For example, you should not use the same name for different printers. If a printer with the same name exists, the new configuration will replace the old one. If a printer cluster with the same name exists, an error appears.
- **Postscript printer**—Check box that allows you to specify if the printer is a postscript-capable printer or not (default is nonpostscript). For postscript printers, the generic.ppd is used. For nonpostscript printers, the PPD is set to raw. Postscript printer configuration supports printing test page and banner pages. This check box is visible only when adding a new printer. If you are modifying a printer, you will not see this check box.
- **Location**—User-specified location of the printer (maximum 127 characters).
- **Description**—User-specified description of the printer (maximum 127 characters).
- **Device URI** (required)—The device URI is the address of the printer, in the following form:

  protocol://server:port/queue

  A maximum of 1024 characters are allowed. The only allowed protocols are lpd, socket, ipp, and http. Validation check on the protocol will be performed but there is no other check on the URI string. A list of example URIs appears to help you enter it properly. You can get the information (printer port, protocol) for your specific printer from the printer manufacturer's manual from the test page of the printer, or from the printer's front panel.

Note Incorrectly specifying the protocol, port, or queue will lead to failure to print pages.

Step 5 Click **Submit** to save your changes.

The new print queue will be visible to clients within one minute.

Step 6 Repeat the process to add other printers to your WAAS print server.

After you add a printer, you should add it to the Windows Active Directory. To add a printer to Active Directory, follow these steps:

Step 1 On the Windows domain controller, open Active Directory Users and Computers.

Step 2 Right-click the container object folder in which you want to publish the printer, click **New**, and then click **Printer**.

The New Object-Printer dialog box appears.

Step 3 In the text box, type the path to the printer, such as \\printserver\printername, and then click **OK**.

Adding Printer Clusters

Printer clustering allows you to group several printers together. This clustering provides failover and load-balancing capabilities. You can include up to 12 printers in a printer cluster. Only printers of the same model and capabilities should be grouped together. Printers in a cluster should be geographically close to each other.

The printer cluster appears as a single print queue to Microsoft clients using the WAAS print server. The WAAS print server does not support the selection of a default printer within a cluster. When you send a print job to a printer in a printer cluster, the print job is automatically forwarded by the WAAS print server to the first available printer.

The default spooling space for all printers (including the printer cluster) is 1 GB.

To add a print cluster, follow these steps:

Step 1 Open the Print Services Administration GUI by doing either of the following:

- From the WAAS Central Manager GUI navigation pane, choose **Configure > Legacy Services > Print Servers**, click the **Edit** icon next to the name of the print server that you want to add to a cluster, and then click the **CUPS** icon in the taskbar.
- From the WAE Device Manager GUI, choose **Configuration** from the Cisco WAE menu, click the **Print Services** tab, and then click the **Open** link.

The Print Services Administration GUI opens.

Step 2 Click the **Printer Cluster** tab.

If you were not previously logged into the Print Services Administration GUI, you are prompted for the user name and password of the print admin account. Enter this information and then click **OK**.

The Printer Clusters window appears. (See Figure 13-5.)

Figure 13-5 Print Services Administration Window—Printer Clusters Tab

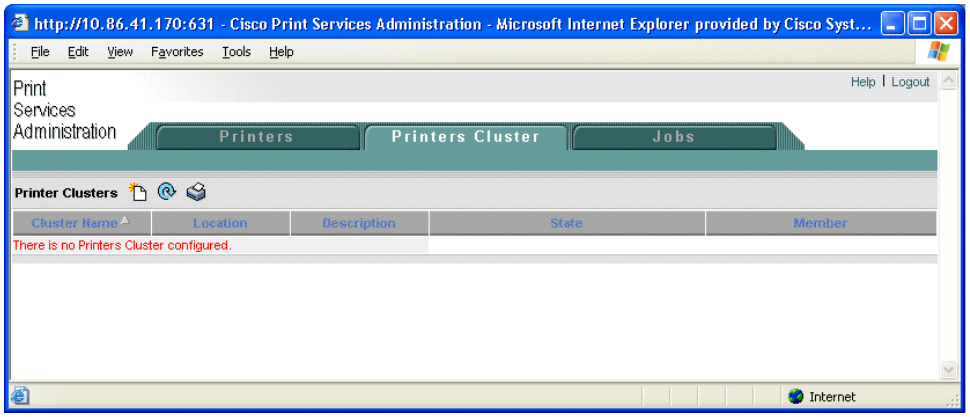

Step 3 Click the **Add Printer Cluster** icon in the taskbar.

The Add Printer Cluster window appears. (See Figure 13-6.)

Figure 13-6 Adding New Printer Cluster

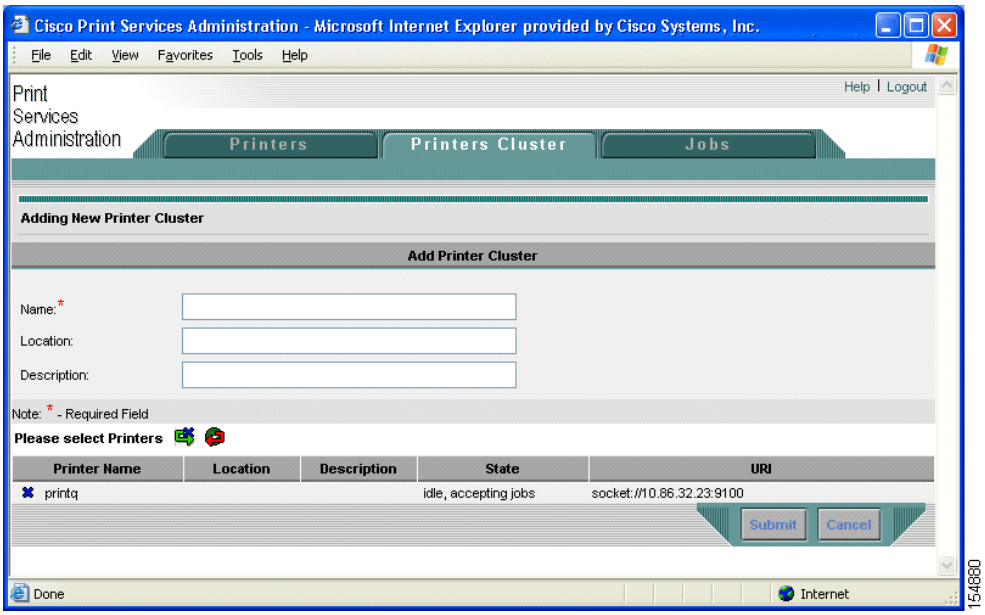

Step 4 Enter a name for the printer cluster.

The printer cluster name can contain letters and numbers, but the only special character that is supported is the underscore (no spaces allowed). The name cannot exceed 127 characters and must be unique across the system. For example, you must not use the same name for different printer clusters. If a printer cluster exists with the same name, the new configuration replaces the old one. If a printer with the same name exists, an error appears.

Step 5 Enter a name for the print cluster location (optional).

The location name cannot exceed 127 characters.

Step 6 Enter a description for the Printer cluster (optional).

This description cannot exceed 127 characters.

Step 7 Choose the printers that you want to join this new cluster by doing one of the following:

- Click to choose all the available printers.
- Click next to each printer you want to join the cluster.

Step 8 Click **Submit**.

The configuration is saved.

The new print queue will be visible to clients within one minute.

Setting Up the WAAS Central Manager as the Driver Repository

You can set up the WAAS Central Manager as a print driver repository that stores all the drivers you want to distribute to your WAAS print servers. Steps 1 through 4 in this section describe how to configure the WAAS Central Manager as the driver repository; steps 5 through 10 describe how to add drivers to the central repository using the Windows Add Printer Driver Wizard.

If you do not set up the WAAS Central Manager as a driver repository, you cannot centrally distribute print drivers to your WAAS print servers. In this case, you must manually install the drivers on each WAAS print server. For more information on installing drivers on an individual print server, see the "Installing Print Drivers on Individual WAAS Print Servers" section on page 13-18.

We recommend that you configure the WAAS Central Manager as a driver repository if you have more than one WAAS print server.

To configure the WAAS Central Manager as the print driver repository, follow these steps:

Step 1 Log into the WAAS Central Manager GUI using the print admin account that you created in the "Creating Accounts with Print Admin Privileges" section on page 13-10.

Step 2 From the WAAS Central Manager GUI navigation pane, choose **Configure > Legacy Services > Print Repository**.

The Print Driver Repository Settings window appears.

Step 3 Check the **Enable Central Manager as Driver Repository** check box.

Step 4 Click **Submit** to save your changes.

Step 5 Log into a Windows client and enter the following command to add drivers to the new central repository:

net use *CM_netbios_name*\print$ * /USER:*username*

The *CM_netbios_name* value is the NetBIOS name of the WAAS Central Manager device, and *username* is the same as the username you created in Step 1.

We recommend that you use postscript (PS) print drivers instead of PCL drivers because postscript drivers load much more quickly.

Step 6 From the Windows client, enter the following command in the Run window:

CM_netbios_name

The *CM_netbios_name* value is the NetBIOS name of the WAAS Central Manager device.

The Central Manager's driver repository appears.

Step 7 Double-click the **Printers** icon.

The Printers on Device window appears.

Step 8 From the Printers on Device window, choose **File > Server Properties**.

The Print Server Properties window appears.

Step 9 Choose the **Drivers** tab, and click **Add...**

The Add Printer Driver Wizard opens.

Note If there is insufficient print share space to add drivers, you will see a Windows error message that indicates that a driver was unable to install because it is not compatible with Windows, or that Windows was unable to save all data due to a failure related to a computer hardware or network connection problem. If you see such a message, remove any drivers that may no longer be in use, then repeat the step to add the new drivers.

Step 10 Proceed through the wizard and add the necessary drivers to the repository on the WAAS Central Manager.

Note To avoid a possible error condition when adding drivers to the repository, do *not* click the Update or Replace button in the Drivers tab of the Print Server Properties window.

Step 11 After completing the wizard, verify that the drivers were successfully added to the repository by choosing **Configure > Legacy Services > Print Drivers** in the WAAS Central Manager GUI.

It may take up to 5 minutes for the list of drivers to appear. Verify that the drivers that you just added appear in this list. You can configure the amount of time it takes the list to be updated by adjusting the System.datafeed.pollRate field described in the "Modifying the Default System Configuration Properties" section on page 9-17.

Note If an error message appears that indicates a driver was unable to install due to a lack of disk space, you will need to free up disk space by first removing any unused drivers from the repository, then repeating the steps to add the driver. Because Windows 95, 98, NT, and 2000 create an extra storage directory during the driver installation process, these operating systems require approximately twice the disk space to install a print driver (requiring that you free up disk space equivalent to at least twice the size of each driver to be installed). Additionally, in Windows 95, 98, and NT operating systems, you can also free up disk space by following the steps below to remove any temporary print driver directories that Windows has not removed. (Windows 2000 and XP automatically remove these temporary directories).

To remove temporary print driver directories, do the following:

a. From the CLI on the WAAS Central Manager, change to the directory where the temporary print driver directories are located using the **cd** command. For example:

   cd spool/samba/printers/W32X86

b. List the directory contents using the **ls** command.

c. Identify the temporary print driver directories by locating directories designated as: "__SKIP__xxx"

d. Use the **rmdir** command to remove the temporary print driver directories.

Installing Print Drivers on Individual WAAS Print Servers

If you do not want to set up the WAAS Central Manager as a driver repository and centrally distribute drivers, you can use the Windows Add Driver Wizard to install the necessary drivers on an individual WAAS print server.

To install printer drivers on a WAAS print server, follow these steps:

Step 1 Log in to the windows client and enter the following command:

net use *WAE_netbios_name*\print$ * /USER:*username*

The *WAE_netbios_name* value is the NetBIOS of the WAAS print server, and *username* is the username you specified in the "Creating Accounts with Print Admin Privileges" section on page 13-10.

Step 2 Using Windows Network Neighborhood, find the Print Server share and click on it. Alternatively, you can open the Run window on the Windows system and enter *wae_netbios_name.*

Step 3 View the printers configured in the WAAS Print Server by opening the Printers folder.

Step 4 Choose the printer for which you want to upload a driver, right click on it, and choose **Properties**. (See Figure 13-7.)

Figure 13-7 Sample Printer Properties Sheet

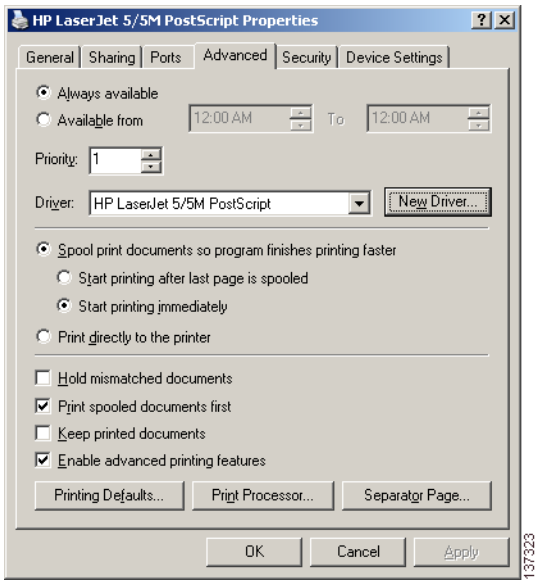

Step 5 Locate the printer driver on the CD-ROM or server.

Step 6 Click the **New Driver** button to start the Add Driver wizard.

Step 7 Click **Next**.

The Printer Driver Selection window appears. (See Figure 13-8.)

Figure 13-8 Driver Selection Window

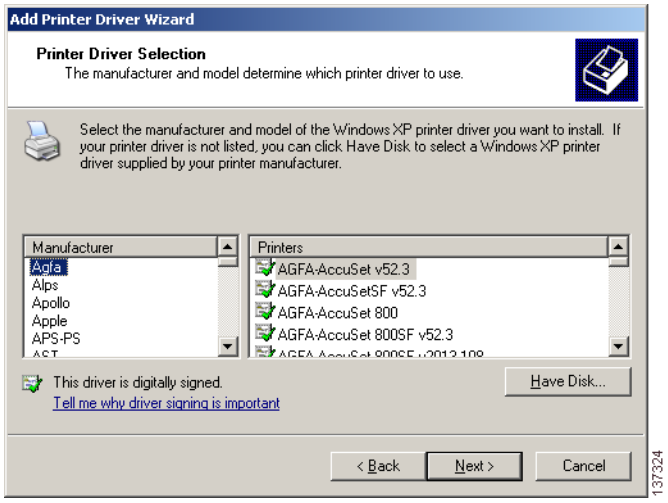

Step 8 If the necessary driver is not displayed in the Printer Driver Selection window, click **Have Disk** and browse for the location of the printer driver on the CD-ROM or server.

Step 9 Click **OK** to finish installing the driver.

The driver has now been uploaded to the WAAS print server so that any user can download the driver to their client machine when printing.

Step 10 Initialize the driver as described in the "Initializing Print Drivers" section on page 13-23.

Distributing Drivers to the WAAS Print Servers

After you have set up the WAAS Central Manager as the driver repository and installed the drivers in the central repository, you can distribute the drivers to your WAAS print servers.

You can distribute drivers to WAAS print servers using one of the following methods:

- Choose the driver first, then choose the print servers to which you want to distribute the driver. This is the recommended method for distributing a single driver to multiple WAAS print servers or device groups. See the "Distributing a Single Driver to Multiple Devices or Groups" section on page 13-20.
- Choose a print server or device group first, and then choose the drivers that should be installed on the selected device. We recommend that you use this method for distributing multiple drivers to a single print server or device group. See the "Distributing Multiple Drivers to a Single Device or Group" section on page 13-21.

After you distribute a print driver to one or more WAAS print servers, you can verify the print driver distributions as described in the "Verifying Print Driver Distribution" section on page 13-21.

Note If you try to distribute a print driver to a WAE that has not been enabled with print services, the driver is not distributed to the WAE.

Distributing a Single Driver to Multiple Devices or Groups

To distribute a single driver to multiple devices or groups, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Configure > Legacy Services > Print Drivers**. The Drivers in the Repository window appears.

Step 2 Click the **Edit** icon next to the driver that you want to distribute. The driver's home window appears.

From this window, you can perform the following tasks:

- Distribute the driver to print servers as described in the steps that follow.
- Delete the driver from the repository by clicking the **Trash** icon in the taskbar.
- View details about the driver.

Step 3 Click one of the following options in the navigation pane to distribute this driver:

- **Distribute to Print Server**—Distributes the driver to an individual WAAS print server.
- **Distribute to Device Group**—Distributes the driver to a group of WAAS print servers that have been included in the same device group.

The Print Server assignments window or the Device Group assignments window appears depending on the selected option.

Step 4 Choose the devices to which you want to distribute the WAAS print driver.

To choose the devices, use one of the following procedures:

- Click in the taskbar to distribute the driver to all available print servers or device groups.
- Click next to each print server or device group to distribute the driver to these specific devices. The icon changes to when selected.

Step 5 Click **Submit**.

The icon next to the selected devices changes to , and the driver is distributed to the specified devices.

Step 6 Verify that the print driver was successfully distributed (see the "Verifying Print Driver Distribution" section on page 13-21).

Distributing Multiple Drivers to a Single Device or Group

To distribute multiple drivers to a single device or group, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group to which you want to distribute the print drivers.

The Device Dashboard window or the Modifying Device Group window appears.

Step 3 From the navigation pane, choose **Configure > Acceleration > Legacy Services > Download Drivers**.

The Drivers in the Repository window appears.

Step 4 Choose the drivers to distribute by doing either of the following:

- Click in the taskbar to choose all the drivers in the list.
- Click next to each driver to distribute. The icon changes to when selected.

Step 5 Click **Submit**.

The icon next to the selected drivers changes to and the selected device or device group downloads the specified drivers.

Step 6 Verify that the print drivers were successfully distributed (see the "Verifying Print Driver Distribution" section on page 13-21).

Verifying Print Driver Distribution

After you distribute a print driver to one or more WAAS print servers, you can use the WAAS Central Manager GUI to check for distribution errors and redistribute the drivers if necessary.

To verify that a print driver has been properly distributed to a WAAS print server, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Configure > Legacy Services > Print Servers**.

The Print Servers table appears. The Driver Distribution Status area of this table shows the number of completed and failed driver distribution attempts. This table also shows the number of drivers that are in the process of being downloaded by the print server.

Step 2 Check the Print Server Errors column and the Failed column for an error count.

These columns show the number of errors reported on a WAAS print server. The Failed column shows the number of driver distribution errors, and the Print Server Errors column shows the number of system errors on the print server.

Step 3 If the Print Server Errors column or the Failed column contains a number other than 0 (for example, 1, 2 or 3), do the following:

a. Click the number in the column. A list of errors appears.

b. Use the error information to resolve the print server or driver distribution problem.

c. Redistribute the failed print drivers by clicking the **Retry Downloading Failed Drivers** icon, and then repeat the previous steps to check for errors.

The Retry Downloading Failed Drivers icon appears in the taskbar on the following pages in the WAAS Central Manager GUI:

- Go to **My WAN > Manage Devices**, click the **Edit** icon next to the appropriate print server, and then choose **Configure > Acceleration > Legacy Services > Drivers** in the navigation pane.
- Go to **My WAN > Manage Device Groups**, click the **Edit** icon for the appropriate device group, and then choose **Configure > Acceleration > Legacy Services > Download Drivers** in the navigation pane.

Step 4 If the Print Server Errors column and the Failed column for your WAAS print servers contains a 0 (no distribution errors), do the following:

a. Click the **Edit** icon next to one of the print servers to which you distributed the print driver.

b. From the navigation pane, choose **Configure > Acceleration > Legacy Services > Print Services > Drivers**.

A list of drivers installed on this device appears.

c. Verify that the Download Status column for the drivers that you just distributed shows Completed.

If the print driver does not appear in the list, wait up to 10 minutes and refresh the page. It can take up to 10 minutes for a newly distributed driver to appear in the list as Completed.

Associating a Driver with your Printer

After you distribute a driver to a WAAS print server, you need to associate the driver with the printer that you added in the "Adding a Printer to the WAAS Print Server" section on page 13-12.

To associate a distributed driver with a printer, follow these steps:

Step 1 Log in to a Windows client.

Step 2 Run the following command to set the print admin privilege level:

net use *WAE_netbios_name*\print$ * /USER:*username*

The *WAE_netbios_name* is the NetBIOS name of the WAAS print server containing the printer you added in the "Adding a Printer to the WAAS Print Server" section on page 13-12, and *username* is the username you specified in the "Creating Accounts with Print Admin Privileges" section on page 13-10.

Step 3 Enter the following command in the Run window on the Windows client:

wae_netbios_name

The printer share window appears.

Step 4 Double-click the printer icon to view a list of printers.

Step 5 Right-click on the printer that you added in the "Adding a Printer to the WAAS Print Server" section on page 13-12 and choose **Properties**.

Step 6 Click the **Advanced** tab, and in the Driver drop-down list, choose the appropriate driver.

Step 7 Click **Apply** to save your changes, and then click **OK** to close the window.

Initializing Print Drivers

Before you install a print driver on a client, you need to initialize the driver so that your branch office clients can successfully install the driver and be able to print from applications such as Microsoft Word and PowerPoint. Uninitialized drivers often do not install successfully and can cause errors when clients use them to print from Microsoft Word and PowerPoint.

Manual initialization of a driver is required because the WAAS print server cannot execute Windows driver code locally to automatically initialize the driver.

To manually initialize a print driver, follow these steps:

Step 1 From a Windows client, log in to the WAAS print server as an administrative user.

Step 2 Open the Printers and Faxes folder on the WAAS print server.

Step 3 Choose the shared printer.

Step 4 Right-click on the printer and choose **Properties**.

Step 5 From the General tab, choose **Printing Preferences**.

Step 6 Change the page orientation from Portrait to Landscape, click **Apply**, and then click **OK**.

Step 7 Reset the page orientation from Landscape back to Portrait, click **Apply**, and then click **OK**.

Changing the page orientation ensures that the setting is applied correctly.

Step 8 (Optional) Set the desired printing defaults, which are then applied to all future client driver installations.

After you complete driver initialization, follow these additional steps before continuing:

Step 1 Run the print_diff utility for each print driver, to verify the driver installation on the Samba print server and to help you understand what printer features may not be supported. For details on this utility, see the Print_Diff Utility section in the document *Using the Print Utilities to Troubleshoot and Fix Samba Driver Installation Problems*.

Step 2 Run the print_fix utility for each print driver to resolve any differences between the Samba and Windows print driver properties. For details on this utility, see the Print_Fix Utility section in the document *Using the Print Utilities to Troubleshoot and Fix Samba Driver Installation Problems*.

Step 3 Run the print_diff utility again for each print driver to verify the driver properties after they were updated by the print_fix utility.

Adding the WAAS Print Server to Your Branch Office Clients

Windows clients who will print using the WAAS print server and previously created print queues must run the Microsoft Add Printer Wizard.

To add a printer to a Windows client, follow these steps:

Step 1 From a Windows client, choose **Printers and Faxes** from the Start menu.

Step 2 Click **Add a Printer**.

The Add Printer Wizard opens.

Step 3 Click **Next**.

Step 4 Click the radio button to choose a network printer.

Step 5 Click **Next**.

The Locating Your Printer window appears. (See Figure 13-9.)

Figure 13-9 Locating Your Printer

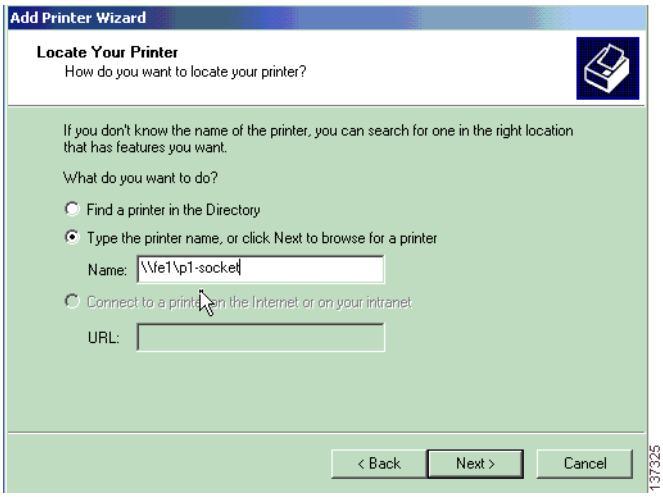

Step 6 Use one of the following methods to choose the printer, and then click **Next**:

- Enter the printer name.
- Find the WAAS print server in the Active Directory (if the printer is published).
- Browse the domain for the WAAS print server.

Note Interoperability issues can occur when Windows browses for the WAAS Print Queue in the Add Printer Wizard. If they occur, you must explicitly type the printer name.

The printer appears in your list of printers. (See Figure 13-10.)

Figure 13-10 Printer Added Successfully

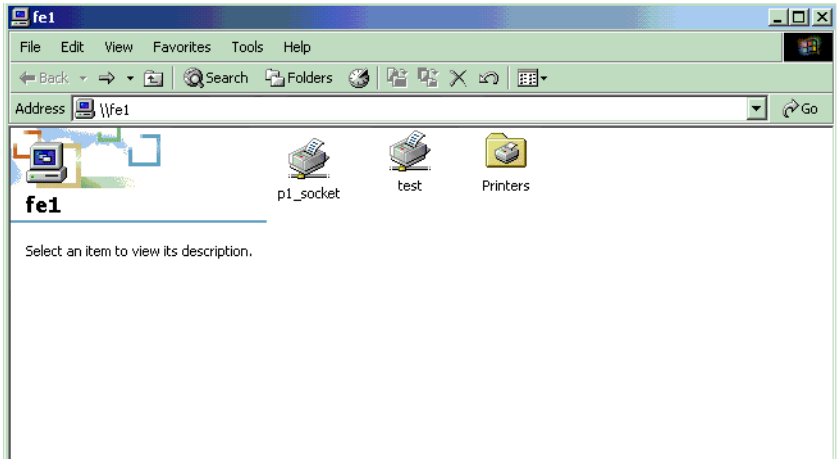

Once the printer is successfully added, you can begin printing.

Managing Print Services

This section contains the following topics:

- Viewing Print Server Details, page 13-25
- Configuring Aggregate Settings, page 13-27
- Using the Print Services Administration GUI, page 13-28
- Viewing the Print Server Page Log, page 13-33

Viewing Print Server Details

Once you set up print services, you can use the WAAS Central Manager GUI to view details about print drivers and the WAAS print servers installed in your network.

Table 13-6 describes how to view print driver and print server details from the WAAS Central Manager GUI.

Table 13-6 Viewing Print Server Details from the WAAS Central Manager GUI

To view	Go to
All drivers that reside in the repository on the WAAS Central Manager	**Configure > Legacy Services > Print Drivers**
Detailed information about a driver	**Configure > Legacy Services > Print Drivers**, and click on the driver.

Table 13-6 Viewing Print Server Details from the WAAS Central Manager GUI (continued)

To view	Go to
Drivers installed on a specific WAAS print server	**Configure > Legacy Services > Print Servers**, and follow these steps: 1. Choose the WAAS print server. 2. Click the **Driver** tab. A list of all the drivers distributed to this WAAS print server appears. For information on how to distribute drivers, see the "Distributing Drivers to the WAAS Print Servers" section on page 13-20.
Printers associated with a WAAS print server	**Configure > Legacy Services > Print Servers**, and follow these steps: 1. Choose the WAAS print server. 2. Click the **Printer** tab (visible by default). Or, go to **My WAN > Manage Devices**, and follow these steps: 1. Choose the WAAS print server for which you want to view the print queue. 2. From the navigation pane, choose **Configure > Acceleration > Legacy Services > Printers**.
A list of WAAS print servers in your network	**Configure > Legacy Services > Print Servers**
A list of WAAS print servers in a specific device group	**My WAN > Manage Device Groups**, then follow these steps: 1. Click the **Edit** icon next to the device group that contains the print servers you want to view. 2. From the navigation pane, choose **Configure > Acceleration > Legacy Services > Print Servers**. A list of print servers associated with the selected device group appears.
Driver distribution errors	**Configure > Legacy Services > Print Servers**, and then click on the number in the Failed column.

Table 13-6 Viewing Print Server Details from the WAAS Central Manager GUI (continued)

To view	Go to
Print server errors	**Configure > Legacy Services > Print Servers**, and then click on the number in the Print Server Errors column.
The status of your WAAS print server	**Configure > Legacy Services > Print Servers**, and then follow these steps: 1. View the Time Stamp column, and ensure that the time is not more than 5 minutes old. Each WAAS print server sends a status report to the WAAS Central Manager on a regular interval specified by the System.monitoring.collectRate field (described in the "Modifying the Default System Configuration Properties" section on page 9-17). If a WAAS print server fails to send this report on time, the server is probably down. 2. If a WAAS print server has not sent a status report in the last 5 minutes, do the following: a. Go to **Admin > Logs > System Messages**. b. Search the message log for the reason the WAAS print server is not functioning correctly. You can use the Node Name column to quickly locate information for a specific WAAS print server. **Note** CUPS log files are rotated when the log file reaches the maximum size of 1MB.

Configuring Aggregate Settings

Aggregate settings allow you to configure a WAAS print server to automatically download drivers that are distributed to a device group to which it is a member. For example, if a WAAS print server belongs to device group DG1 and is configured with aggregate settings enabled, the print server will automatically download any driver that you distribute to the DG1 device group. However, if the same print server is configured with aggregate settings disabled, the print server will not download drivers that you distribute to the DG1 device group.

The WAAS Central Manager does not redistribute drivers that already exist on a WAE device. For example, if a WAAS print server belongs to two device groups (DG1 and DG2) and the aggregate setting is enabled, the print server only downloads one instance of the driver in the event it is distributed to both device groups.

When you add a WAAS print server to a device group, the aggregate settings for the print server are enabled and the drivers belonging to the device group are automatically distributed to the print server. When you remove a print server from a device group when aggregate settings are enabled, all the drivers assigned to the device group are removed from the print server if the drivers are not being used by any printer.

To configure aggregate settings on a WAAS print server, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

The Devices window appears.

Step 2 Click the **Edit** icon next to the WAAS print server for which you want to configure aggregate settings.

The Device Dashboard window appears.

Step 3 From the navigation pane, choose **Configure > Acceleration > Legacy Services > Drivers**.

A list of drivers installed on the WAAS print server appears.

Step 4 Click the **Yes** radio button at the top of the list to enable aggregate settings, or click **No** to disable aggregate settings.

A dialog box asks if you want to continue.

Step 5 Click **OK**.

If you enabled aggregate settings, the WAAS print server begins downloading all the drivers that have been distributed to its device group.

Using the Print Services Administration GUI

The Print Services Administration GUI allows you to perform a variety of tasks for a specific WAAS print server. You can open the Print Services Administration GUI from the WAAS Central Manager GUI or from the WAE Device Manager GUI.

Note After accessing the Print Services Administration GUI, you are prompted for the user name and password of the print admin account when you perform any task such as adding a printer. For information on creating a print admin account on a WAE device, see the "Creating a New Account" section on page 7-4.

This section contains the following topics:

- Opening the Print Services Administration GUI, page 13-28
- Adding a Printer, page 13-29
- Modifying the Printer Configuration, page 13-29
- Enabling Print Banners, page 13-31
- Setting Up Print Clusters, page 13-32
- Viewing Print Jobs, page 13-32

Opening the Print Services Administration GUI

You can open the Print Services Administration GUI from either the WAE Device Manager GUI or from the WAAS Central Manager GUI.

To open the Print Services Administration GUI from the WAE Device Manager GUI, follow these steps:

Step 1 From the WAE Device Manager GUI, choose **Cisco WAE > Configuration**.

Step 2 Click the Print Services tab, and click the **Open** link. The Print Services Administration GUI opens.

To open the Print Services Administration GUI from the WAAS Central Manager GUI, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Configure > Legacy Services > Print Servers**.

A list of the WAAS print servers installed in your network appears.

Step 2 Click the **Edit** icon next to the WAAS print server that you want to manage.

Step 3 Click the **CUPS** icon in the tool bar.

The Print Services Administration GUI appears.

Adding a Printer

For instructions on how to add a printer to a WAAS print server, see the "Adding a Printer to the WAAS Print Server" section on page 13-12.

Modifying the Printer Configuration

To modify the configuration of a an existing printer, follow these steps:

Step 1 From the Printers tab in the Print Services Administration GUI, choose the printer to modify.

The Modifying Printer window appears. (See Figure 13-11.)

Figure 13-11 Modifying Printer Configuration

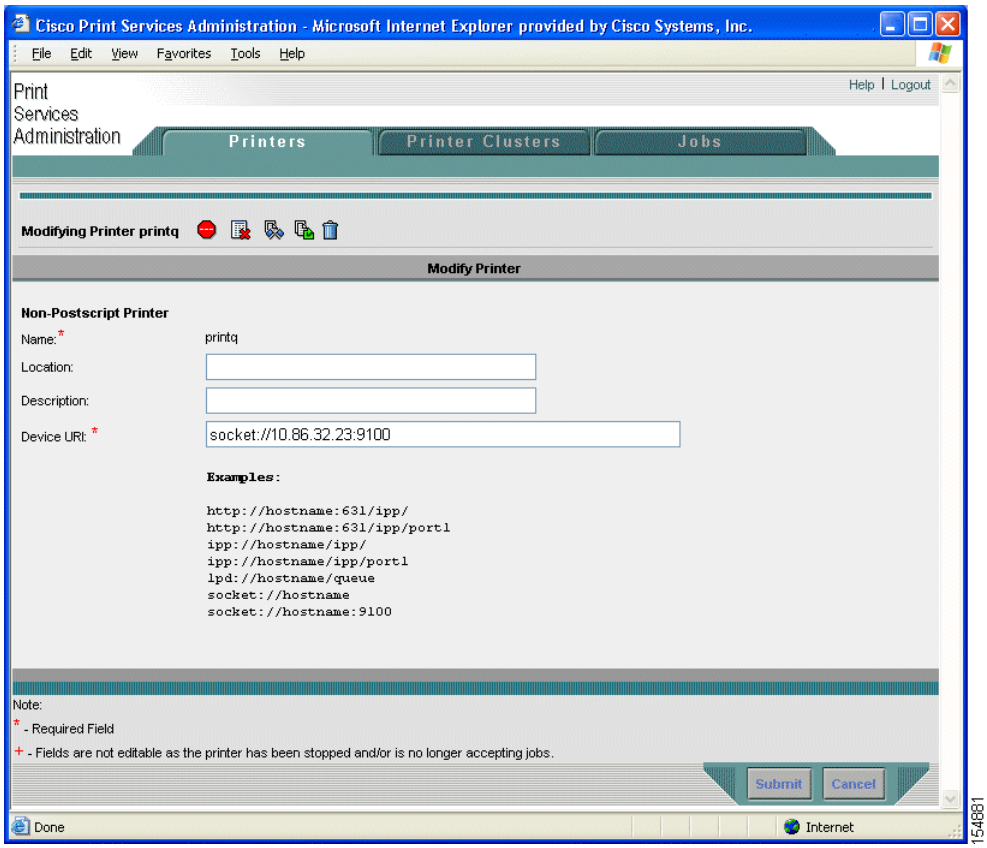

Use this window to modify the location, description, and device URI settings. The name field cannot be modified. You can also perform the following tasks by clicking the corresponding icons:

- Print a test page (for Postscript printer only)—Sends a test page to the printer.
- Configure printer (for Postscript printer only)—Displays the Configure Printer window that can be used to set starting and ending banner page options. For more information, see the "Enabling Print Banners" section on page 13-31.
- Stop/Start the printer—Allows you to stop and start the printer. A warning appears to ask for user confirmation before stopping the printer.
- Accept/Reject Jobs—Allows you to determine if any job can be sent to this printer.
- Show Completed Jobs—Displays the Job Listing Page with only completed jobs for this printer. For more information, see the "Viewing Print Jobs" section on page 13-32.
- Show Active Jobs—Displays the Job Listing Page with only active jobs for this printer. For more information, see the "Viewing Print Jobs" section on page 13-32.
- Delete this printer configuration—Allows you to delete this printer. If the printer belongs to a printer cluster, deleting the printer will remove that printer from the cluster. If the printer is the only printer in a cluster, the cluster itself is removed.

Step 2 Modify any of the following configuration settings:

- Location—User-specified location of the printer (maximum 127 characters).
- Description—User-specified description of the printer (maximum 127 characters).

- Device URI (required)—The device URI is the address of the printer, in the form:

  protocol://server:port/queue

  A maximum of 1024 characters are allowed. The only allowed protocols are lpd, socket, ipp, and http. Validation check on the protocol will be performed but there is no other check on the URI string. A list of example URIs appears to help you enter it properly. You can obtain the information (printer port, protocol) for your specific printer from the printer manufacturer's manual, from the test page of the printer, or from the printer's front panel.

Note You can configure these fields only if the printer is running. If the printer is stopped or is rejecting jobs, the fields are read-only and a message appears indicating the printer status.

Step 3 Click **Submit** to save your changes.

The print queue change will be visible to clients within 1 minute.

Enabling Print Banners

To enable a starting or ending banner for each print job, follow these steps:

Step 1 From the Printers tab in the Print Services Administration GUI, choose the printer to modify.

The Modifying Printer window appears. (See Figure 13-11 on page 13-30.)

Step 2 Click the **Configure** button ( ) in the tool bar.

The Configuring Printer window appears. (See Figure 13-12.)

Figure 13-12 Configuring the Starting and Ending Print Banners

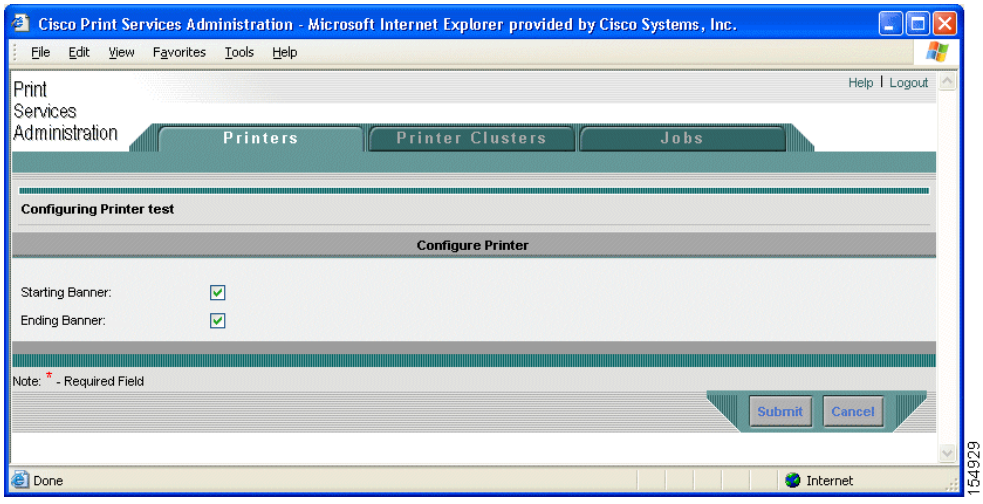

Step 3 Check the check box next to the banner option that you want to enable:

- **Starting Banner**—Enables printing of a starting banner for each job.
- **Ending Banner**—Enables printing of an ending banner for each job.

The starting and ending banners print the same information. These options just allow you to choose the location of the banner on the page.

Step 4 Click **Submit**.

Note If you incorrectly specify a nonpostscript printer as a postscript printer and enable the banner page, the printer will print something similar to "%!PS-Adobe-3.0. %% …" then numerous blank pages. This message is an indication that the printer setup is incorrect. To print valid test page and banner pages, you must specify the printer postscript capability correctly.

Setting Up Print Clusters

For information on setting up print clusters, see the "Adding Printer Clusters" section on page 13-14.

Viewing Print Jobs

Use the Jobs tab in the Print Services Administration GUI to view a list of completed and active jobs.

The pages show the following job details:

- ID —Generated by concatenating the printer name with a sequence number.
- Name—Name of the job you selected to print.
- Size—Size of the job.
- Status—Reported status of the job; only available in Active Job Listing Page.
- Completion Time—Reported completion time of the job; only available in Completed Job Listing window.

The Completed Job Listing window (see Figure 13-13) only shows jobs that are still in the spool area (/local/local1/spool/cups). The spool area retains the last 500 completed jobs and has a 1-GB size limitation. Once the completed job limit is reached, the oldest job is removed to make room for the latest completed job. When a job is removed from the spool area, it is no longer listed in the Completed Job Listing window.

Figure 13-13 Completed Job Listing Window

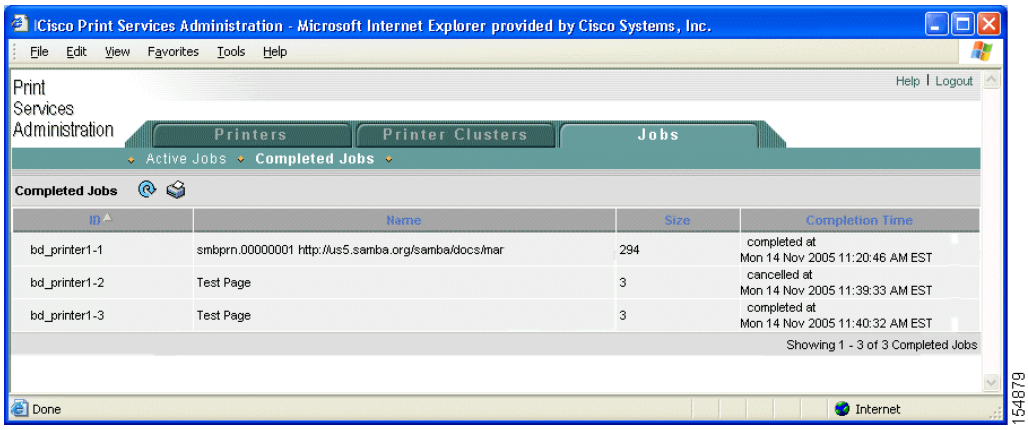

In the Active Job Listing window (see Figure 13-14), you can also perform the following tasks by clicking the corresponding icons next to the job:

- Stop Job—Stops the selected job.
- Delete Job—Deletes the selected job.

The Active Job Listing window shows all the print jobs that are not completed or cancelled.

Figure 13-14 Active Job Listing Window—Job Can Be Stopped or Deleted

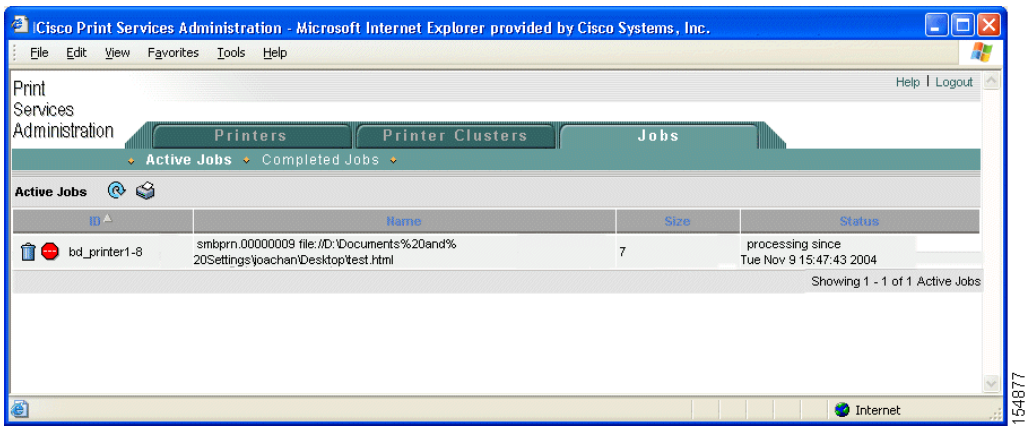

Viewing the Print Server Page Log

The WAAS print server page log is stored in /local/local1/logs/cups_access_log. The cups_access_log file lists each page that is sent to a printer. Each line contains the following information:

printer user job-id date-time page-number num-copies job-billing hostname

Here is an example of a log entry:

DeskJet root 2 [20/May/2008:19:21:05 +0000] 1 0 acme-123 localhost

The field descriptions are as follows:

- The *printer* field contains the name of the printer that printed the page. If you send a job to a printer class, this field will contain the name of the printer that was assigned the job.
- The *user* field contains the name of the user (the IPP requesting-user-name attribute) that submitted this file for printing.
- The *job-id* field contains the job number of the page being printed. Job numbers are reset to 1 whenever the CUPS server is started, so do not depend on this number being unique.
- The *date-time* field contains the date and time when the page started printing. The format of this field is identical to the *date-time* field in the access_log file.
- The *page-number* and *num-pages* fields contain the page number and number of copies being printed of that page. For printers that cannot produce copies on their own, the *num-pages* field will always be 1.
- The *job-billing* field contains a copy of the job-billing attribute provided with the IPP create-job or print-job requests or "-" if none was provided.
- The *hostname* field contains the name of the host (the IPP job-originating-host-name attribute) that originated the print job.

Troubleshooting Print Services

Troubleshooting issues are documented in the following categories:

- General Known Issues, page 13-34
- Login and Access Problems, page 13-35
- Avoiding Print Problems, page 13-35
- Understanding Interactions Between the WAAS Central Manager and the WAAS CLI, page 13-36

Also see the document *Using the Print Utilities to Troubleshoot and Fix Samba Driver Installation Problems*, which explains how to troubleshoot and fix many common Samba print driver installation problems using the print utilities tools.

General Known Issues

The following are general known issues with WAAS print services:

- Linksys devices do not support IPP 1.1 printing; use the LPD protocol only.
- A Linksys URI should use L1/2/3 instead of P1/2/3 as specified in the Linksys user guide.
- In some cases the Windows Print Queue status may not refresh correctly on a Windows client. If this happens, press F5 to manually refresh.
- After adding a printer in the Print Services Administration GUI, it can take up to one minute for the printer to appear on a Window client system.
- In some cases the printer status and error messages are not displayed to the print client.
- The only printers that are supported are IP-based network printers. If you have directly attached printers, use Parallel/Serial/USB IP adapters.
- If you receive an “Access is denied” error when trying to load a driver to the central repository, ensure that the driver file is marked as Read/Write and not Read Only.

- Print problems can be caused by corrupted TDB files. To verify these files, use the **windows-domain diagnostics** CLI command. See the *Cisco Wide Area Application Services Command Reference* for more information about this command.

Login and Access Problems

The following login and access issues can occur when using WAAS print services:

- Authentication information is contained in the **syslog.txt** file. You can raise the log level in this file by adding the relevant smb.conf directives using the **smb-conf** advanced command. Refer to the *Cisco Wide Area Application Services Command Reference* for more information.
- If you cannot access the WAAS print server from Windows, you should log out of Windows and enter the following command:

  net use *WAE_netbios_name*\print$ * /USER:*username*

  The *WAE_netbios_name* value is the NetBIOS name of the WAAS Central Manager device, and *username* is the username you specified in the "Creating Accounts with Print Admin Privileges" section on page 13-10.
- If the Admin user cannot add a driver to a printer, the user should log out of Windows and enter the following command:

  net use *WAE_name*\print$ * /USER:*username*

  The *WAE_name* is the hostname or IP address of the WAAS Central Manager device, and *username* is the name of the account created in the "Installing Print Drivers on Individual WAAS Print Servers" section on page 13-18.
- To avoid issues with usernames, follow these guidelines:
  - The WAAS Central Manager checks for unique usernames in a case-insensitive manner (for example, "user1" is identical to "User1," and you cannot add both).
  - The CLI command for creating users is case-sensitive, you can add both user1 and User1 using the CLI.

Avoiding Print Problems

To avoid printing problems when using WAAS print services, follow these guidelines:

- When using a cluster, always use the cluster configuration window and do not configure the cluster members individually. This action may lead to inconsistent behavior, even when it is a cluster of one.
- Printer clusters should contain printers with similar capabilities. Failure to configure correctly may cause printing error. Deletion of a printer will also remove that printer from its associated clusters.
- Print test page and banner pages are only supported by PostScript printers. Nonpostscript printers will print unreadable output.
- If Point and Print does not work, ensure that the drivers have been properly distributed to the WAAS print servers as described in the "Distributing Drivers to the WAAS Print Servers" section on page 13-20.
- Jobs are reprinted from the beginning. If WAAS print services are restarted due to configuration changes, in-process jobs will be reprinted.

- If jobs are not being printed, do the following:
  - Ensure that the printer is functional.
  - Ensure that the URI is set correctly.
  - Ensure that the printer is not stopped or rejecting jobs in the WAAS print services configuration.
  - Ensure that either the CIFS accelerator or the legacy WAFS edge services on the print server are enabled and running.
  - Check if both **smbd** and **cupsd** processes are running.
  - Check if the spool directories are full.
  - Ensure that the printer driver has been properly distributed. For more information, see the "Initializing Print Drivers" section on page 13-23.
- Minimize changes to the NetBIOS name of the WAAS print server. If you need to change the NetBIOS name, make sure to the use either the WAAS Central Manager GUI or the **windows-domain** CLI command.
- If a PostScript job is sent to a nonpostscript printer, you will see a "%!PS-Adobe-3.0..." message. To fix this, configure the printer as nonpostscript, or send the print job to a PostScript printer.
- If you cannot check the check box for advanced features on a printer in the branch, it may be that "Enable advanced printing features" is allowing the client printer driver to communicate with the print server using Enhanced Metafile Spooling (EMF).

  WAAS print services supports only RAW spooling and not EMF because no rendering is done on the WAAS print server. The client renders the print job, and then submits it in a RAW format, as if the printer is going to receive that job.

Note No printer functionality is lost because of RAW spooling. The only functionality lost is the print rendering by the WAAS print server. In RAW spooling, the client is doing the rendering instead of having it done by the print server itself.

Understanding Interactions Between the WAAS Central Manager and the WAAS CLI

Because you can use the WAAS Central Manager or the WAAS CLI to configure print services, you should be aware of the following interactions between these tools:

Note Restarting CUPS may temporarily disrupt print jobs.

- If Print Services is not running, the Open link for Print Services Administration GUI is disabled in the WAAS Central Manager GUI. If you stop print services by entering the **no print-services enable** global configuration command in the CLI after the WAAS Central Manager is loaded, the message "The page cannot be displayed" appears after attempting to access print services.
- Some configuration changes made from the WAAS Central Manager GUI or the CLI require Samba and CUPS software to be restarted, temporarily interrupting print services availability and operation. User confirmation is required before configuration changes are made.

For more information about the print services CLI commands, see the *Cisco Wide Area Application Services Command Reference*.

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 only on certain models of WAE and WAVE devices. On unsupported WAE and WAVE 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, page 14-1
- Preparing to Use Virtual Blades, page 14-3
- Configuring Virtual Blades, page 14-4
- Enabling and Disabling Virtual Blades, page 14-9
- Copying a Disk Image to a Virtual Blade, page 14-10
- Backing Up and Restoring a Virtual Blade, page 14-12

About Virtual Blades

WAAS virtual blades act as computer emulators within your WAE or WAVE 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 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.

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 WAE device, and the rest of your WAAS network.

Note When you configure a virtual blade on your WAE or WAVE 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 WAE.

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 14-9.
4. Transfer files to the virtual blade.	Copies files to the WAE 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 14-12.

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 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the WAE device that you want to configure.

Note You can enable virtual blades only on application accelerator WAEs, not on Central Manager devices.

Note You can configure virtual blades only on individual WAAS devices. You cannot configure virtual blades on device groups.

Step 3 From the navigation pane, 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 "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 WAE or WAVE device, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the WAE device that you want to configure.

Step 3 From the navigation pane, choose **Admin > Virtualization > Virtual Blades**. The Virtual Blade Entries window appears (see Figure 14-1).

Figure 14-1 Virtual Blade Entries List Window

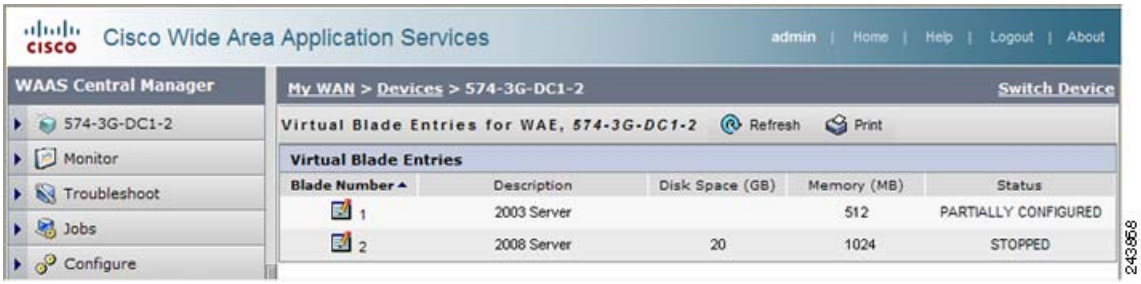

Any existing virtual blades are displayed in the Virtual Blade Entries list.

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 14-4).

Step 4 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

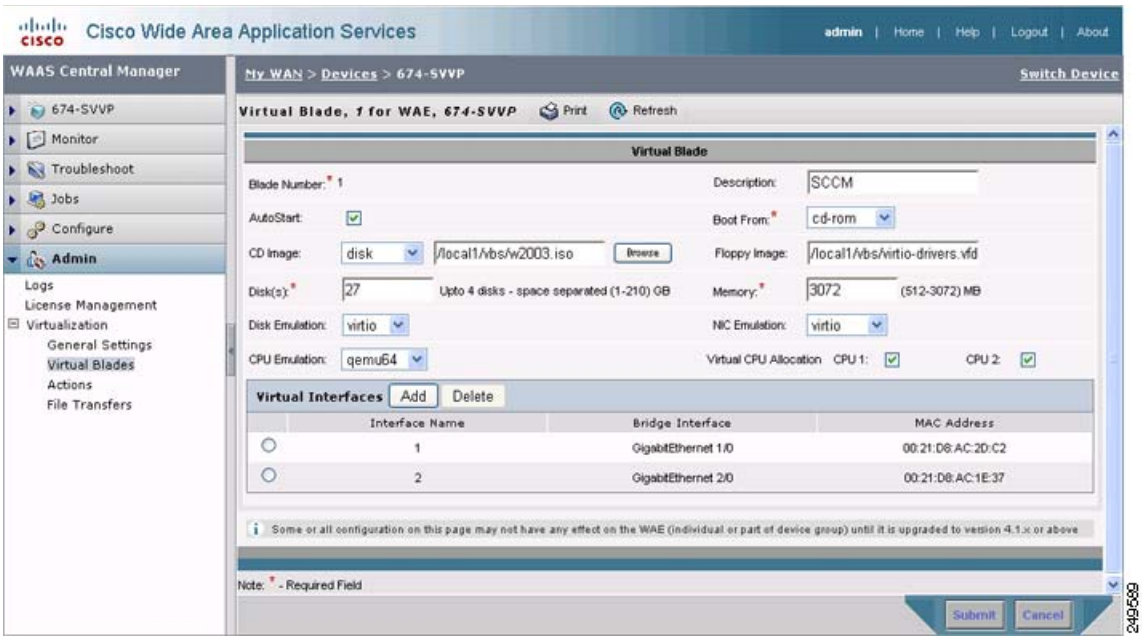

Step 5 Configure the virtual blade system parameters as needed to run your operating system and applications:

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 WAE 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 WAE CD-ROM drive.
- Choose **disk** to read the CD image from an ISO file on the WAE 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 WAE device that is running WAAS version 4.1.1, if multiple virtual hard disks are configured on the WAE device. If you use this Central Manager window to change any part of the virtual blade configuration of a WAE device running WAAS version 4.1.1 that already has multiple virtual hard disks configured, the Central Manager will remove the disk configuration for all disks after the first disk, and erase the other virtual disks.

You cannot configure multiple virtual hard disks from the Central Manager for WAE 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 that is assigned to the virtual blade.

If a single CPU is chosen, only that CPU is used. If the CPU list contains two entries, both CPUs are used in SMP mode. For more than one CPU, odd numbered virtual blades use CPU 1, and even numbered virtual blades use CPU 2.

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. On a device with two CPUs, one CPU is always reserved for the WAAS software. On a device with four CPUs, two are reserved for the WAAS software. If no virtual blades are started, all CPUs are used for the WAAS software.

Step 6 Configure the interface bridge you want to use between the virtual blade and the physical interfaces on your WAE by doing the following:

a. In the Virtual Interfaces pane, click **Add**. The Virtual Interface Add pane is displayed (see Figure 14-3).

Figure 14-3 Virtual Interfaces Add Display Pane

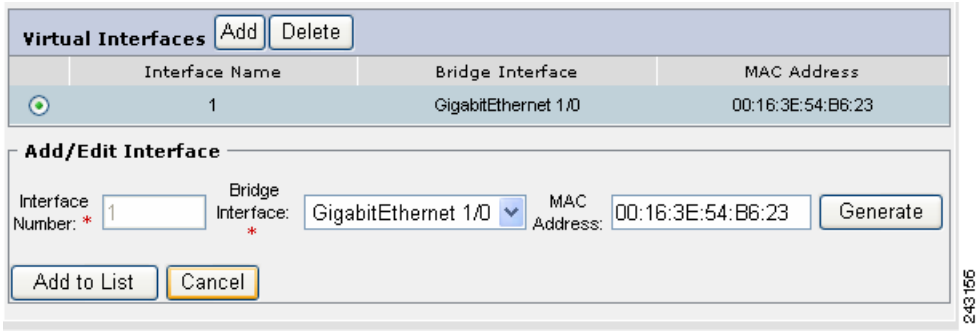

b. In the **Interface Number** field Enter the virtual blade interface to be bridged. Valid values are 1 or 2.

c. In the **Bridge Interface** list, choose the physical WAE interface that the virtual blade interface is bridged to. Choose either **GigabitEthernet** or **PortChannel**.

d. In the MAC Address field, enter the MAC address of the bridged interface or click **Generate** to have WAAS generate the MAC address for you.

e. Click **Add to List** to add the virtual interface to the virtual blade interface list.

Note To access the virtual blade console use the IP address of the bridge interface, with the virtual blade number specified as the port number (separated by a colon). For example if you bridged the GigabitEthernet 1/0 interface, and its IP address is 10.10.10.20, use **10.10.10.20:1** to get to the virtual blade 1 console.

Step 7 Choose the virtual interface by clicking its radio button on the display.

Step 8 Click **Submit**.

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 an interface on your WAE
- **(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 Software Center site:

http://www.cisco.com/public/sw-center/index.shtml

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 GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the WAE device that you want to configure.

Step 3 From the navigation pane, choose **Admin > Virtualization > Actions**. The Virtual Blade Actions window appears (see Figure 14-4).

Figure 14-4 Virtual Blade Actions Window

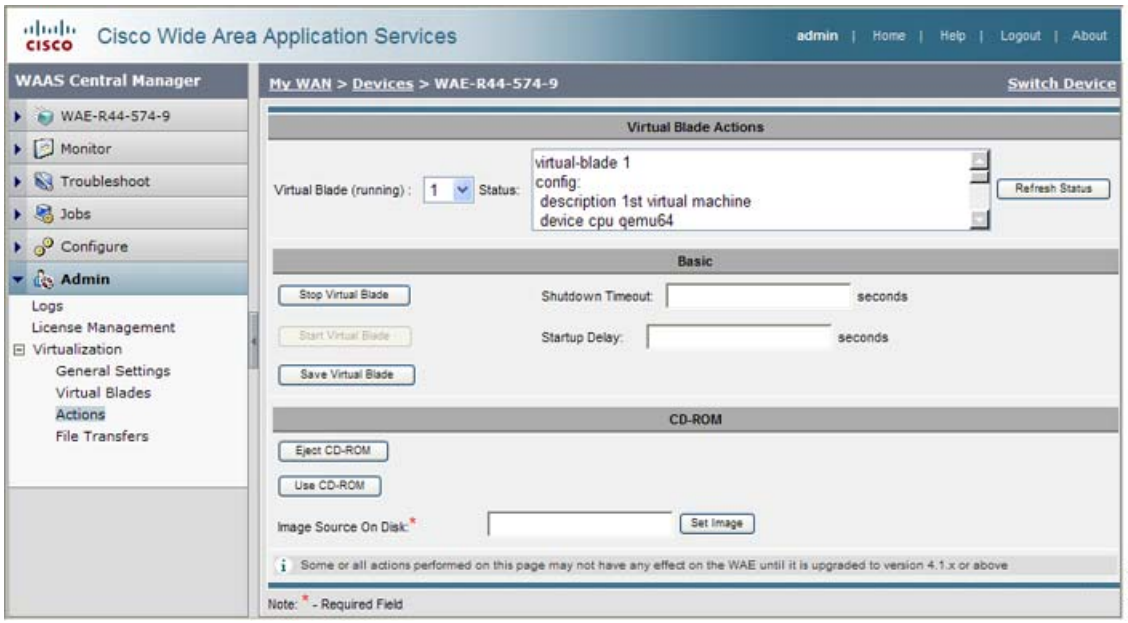

Step 4 In the Virtual Blade list, choose the virtual blade that you want to enable or disable. The status of the virtual blade is displayed in the Status field.

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 5 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 6 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 7 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 GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the WAE device that you want to configure.

Step 3 From the navigation pane, choose **Admin > Virtualization > File Transfers**. The Virtual Blade File Transfers window appears (see Figure 14-5).

Figure 14-5 Virtual Blade File Transfers Window

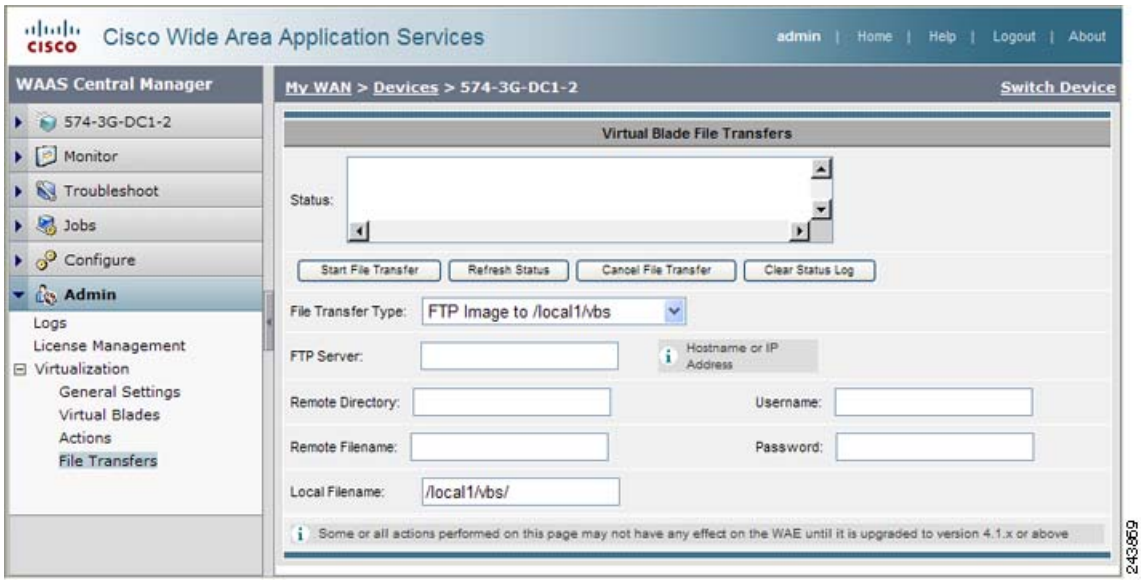

Step 4 In the File Transfer Type list, choose **FTP Image to /local1/vbs** (the default).

Step 5 In the FTP Server field, enter the IP address or hostname of the FTP server where the disk image resides.

Step 6 In the Remote Directory field, enter the path to the directory on the FTP server where the disk image resides.

Step 7 In the Remote Filename field, enter the filename of 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 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 10 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 14-12.

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 14-9.

Step 2 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

Step 3 Click the **Edit** icon next to the WAE device that you want to configure.

Step 4 From the navigation pane, choose **Admin > Virtualization > File Transfers**. The Virtual Blade File Transfers window appears (see Figure 14-5).

Step 5 In the File Transfer Type list, choose **Backup Virtual Blade to FTP**.

Step 6 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 7 In the Remote Directory field, enter the path to the directory on the FTP server where you want to copy the disk image.

Step 8 In the Remote Filename field, enter the filename for the disk image.

Step 9 In the Username and Password fields, enter a valid username and password for the FTP server.

Step 10 In the Virtual Blade No. field, enter the number of the virtual blade that you want to back up.

Step 11 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 12 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**.

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
```

PART 4

Maintaining, Monitoring, and Troubleshooting your WAAS Network

CHAPTER 15

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, page 15-1
- Backing Up and Restoring your WAAS System, page 15-9
- Performing Disk Maintenance for RAID-1 Systems, page 15-22
- Replacing Disks in RAID-5 Systems, page 15-23
- Configuring the Central Manager Role, page 15-24
- Enabling Disk Encryption, page 15-28
- Configuring a Disk Error-Handling Method, page 15-29
- Enabling Extended Object Cache, page 15-30
- Activating All Inactive WAAS Devices, page 15-31
- Rebooting a Device or Device Group, page 15-32
- Performing a Controlled Shutdown, page 15-33

Upgrading the WAAS Software

Table 15-1 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. This is a change from WAAS version 4.0.x, where the WAAS Central Manager should be the lowest version. For details on version interoperability limitations, see the *Release Note for Cisco Wide Area Application Services*.

If the WAAS Central Manager (version 4.2.1) 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 4.2.1 can manage WAE devices that are running version 4.0.19 and later. Some WAAS Central Manager pages (with new features) are not applicable to WAAS devices that are running a version lower than 4.2.1. If you modify the configuration on such pages, the configuration is saved, but it has no effect on the device until the device is upgraded to version 4.2.1.

Note WAAS version 4.2.1 does not support running in a mixed version WAAS network where any WAAS device is running a software version lower than 4.0.19. If you have any WAAS devices running version 4.0.17 or earlier, you must first upgrade them to version 4.0.19 (or a later 4.0.x version), before you install version 4.2.1. You should first upgrade any WAEs to version 4.0.19 (or a later 4.0.x version) and then upgrade any WAAS Central Managers to version 4.0.19 (or a later 4.0.x version). Then, you should begin the upgrade to version 4.2.1, starting with the WAAS Central Managers first.

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 15-1 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 15-3.
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 15-3.
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 15-4.
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 15-6.
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 15-8.
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 15-9.

If you need to downgrade or roll back the WAAS software from version 4.2.1 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 **My WAN > Manage Devices**. The Devices window displays the software version for each device listed.

You can also click 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 navigation pane for any given device, choose **Troubleshoot > 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 preferred web browser and open this location:

http://www.cisco.com/cisco/web/download/index.html

Step 2 Click the **Application Networking Services** product category.

Step 3 When prompted, log in to Cisco.com using your designated username and password. The Download Software window appears, listing the available software products.

Step 4 Click the **Wide Area Application Software** folder to expand it.

Step 5 Choose the **Cisco Wide Area Application Services (WAAS) Software** page.

Step 6 Choose a link to the WAAS software version that you want. The Software Download pane appears to the right side of the page.

Step 7 Click the **Download Now** button next to the file that you want to download. Alternatively, click the **Add to cart** button if you want to download multiple files at once, then click the **Download Cart** link to view your cart.

Be sure to specify the correct type of software image for the devices you want to upgrade, either a Universal (WAAS-4.2.*x*.*x*-K9.bin) or Accelerator only (WAAS-4.2.*x*.*x*-K9.AA.bin) image.

The Download Cart appears.

Step 8 Click the **Proceed With Download** button to download the file.

- If this is the first time you have downloaded a software file from Cisco.com, the Encryption Software Export Distribution Authorization form appears.

  - Fill out the form and click **Submit**. The Cisco Systems Inc., Encryption Software Usage Handling and Distribution Policy appears.
  - Read the policy and click **Accept**. The Cisco End User Software License Agreement appears.
- If you previously filled out the Encryption Software Export Distribution Authorization form and read and accepted the Cisco Systems Inc., Encryption Software Usage Handling and Distribution Policy, these forms are not displayed again. Instead, the Cisco End User Software License Agreement appears after you click the **Proceed With Download** button.

Step 9 Read the Cisco End User Software License Agreement and click the **Agree** button. The Download option window appears.

Step 10 Choose either the **Download Manager Option** (uses Java) or the **Non Java Download Option** link.

- If you choose the Download Manager Option, a Java window appears. The first time you use the Download Manager, you will also see a pop-up window with a security warning; click the **Run** button to install the Download Manager. In the Select Location pop-up window, choose the download location and click **Open**. Your download begins and you can use controls in the Download Manager window to pause, resume, or cancel the download.
- If you choose the Non Java Download Option, a pop-up window opens with links to download your file. Right-click the software file link to download the software and use the **Save Link As** or the **Save Link Target As** option to save the file to your FTP or HTTP server.

Step 11 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. The software file settings form in the WAAS Central Manager GUI defines the software file (.bin) and can be used to specify how to obtain the software file, and whether to preposition it or download it directly to a device.

There are two types of WAAS software files, as follows:

- Universal—Includes Central Manager and Application Accelerator functionality. You can use this type of software file to upgrade either a Central Manager or an Application Accelerator device.
- Accelerator only—Includes Application Accelerator functionality only. You can use this type of software file to upgrade only an Application Accelerator device. If you want to change an Application Accelerator 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 GUI navigation pane, choose **Jobs > Software Update**.

Step 2 Click the **Create New Software File** icon in the taskbar.

The Creating New Software File window appears. (See Figure 15-1.)

Figure 15-1 Creating New Software File Window

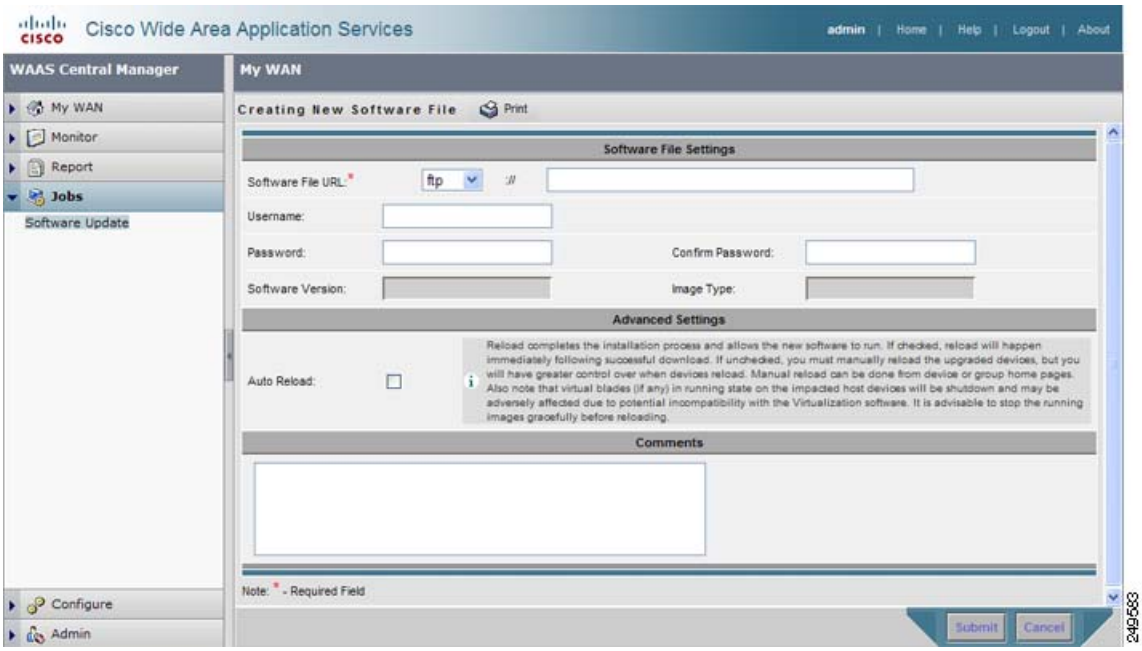

Step 3 In the Software File URL field, specify the location of the new WAAS software file as follows:

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-4.x.x.x-K9*.bin

where *WAAS-4.x.x.x-K9* 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 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, 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.

Note With WAAS version 4.0.x, you should upgrade the WAE devices before the WAAS Central Manager.

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 (version 4.2.1) 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, 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 GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon of the Central Manager device that you want to upgrade.

The Device Dashboard window appears.

Step 3 Verify that the device is not already running the version to which you plan to upgrade.

Step 4 Click the **Jobs > Update Software** button.

The Software Update window appears.

Step 5 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 15-4.

Step 6 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 15-2 for a description of upgrade status messages.

Step 7 If you are using Internet Explorer browser to upgrade the primary WAAS Central Manager, clear the cache before closing the browser and restarting the browser session to the WAAS Central Manager.

You can clear the browser cache in the Tools > Internet Options > General tab of Internet Explorer.

Note For the WAAS Central Manager GUI to work properly, you must first clear the cache in Internet Explorer before restarting 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 15-2 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.

Table 15-2 Upgrade Status Messages (continued)

Upgrade Status Message	Condition
Proceeding with Flash Write	A request has been made to write the software file to the device flash memory.
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 GUI navigation pane, choose **My WAN > Manage Device Groups**.

The Device Groups listing window appears, listing all the device groups in your WAAS network.

Step 2 Click the **Edit** icon next to the name of the device group that you want to upgrade.

The Modifying Device Group window appears.

Step 3 In the navigation pane, choose **Jobs > Software Update**.

The Software Update for Device Group window appears.

Step 4 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 15-4.

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 5 Click **Submit**.

To view the progress of an upgrade, go to the Devices window (**My WAN > Manage 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 15-2 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 GUI navigation pane, choose **Jobs > 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, page 15-10
- Backing Up and Restoring a WAE Device, page 15-11
- Using the Cisco WAAS Software Recovery CD-ROM, page 15-12
- Recovering the System Software, page 15-17
- Recovering a Lost Administrator Password, page 15-19
- Recovering from Missing Disk-Based Software, page 15-20

- Recovering WAAS Device Registration Information, page 15-21

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 The CMS database backup does not back up print drivers. When you perform a Central Manager database backup, you must reinstall your print drivers if you are using WAAS print services.

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:

```
CDM# cms database backup
creating backup file with label 'backup'
backup file local1/cms-db-7-22-2008-17-36_4.1.3.0.1.dump is ready. 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-7-22-2008-17-36_4.1.3.0.1.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:

```
CDM# cd /local1
CDM# 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:

```
CDM# no cms enable
```

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:

```
CDM# cms database delete
```

c. Restore the CMS database contents from the backup file:

```
CDM# 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
```

d. Enable the CMS service on the Central Manager:

```
CDM# 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 15-10.

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 10-7.
- CLI—You can use the following commands to back up and restore a device's database:
  - **wafs backup-config**—Saves the entire legacy WAFS system configuration to a file, including the configuration for file servers, printers, and users. We strongly recommend that you register your WAE again after you use this command.
  - **wafs restore-config**—Restores configuration based on the specified backup file. This command automatically performs a reload function.
  - **copy running-config**—Saves the currently running network configuration to the startup 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*.

Using the Cisco WAAS Software Recovery CD-ROM

A software recovery CD-ROM ships with WAE hardware devices. This section contains instructions for using the software recovery CD-ROM to reinstall your system software if for some reason the software that is installed has failed.

Caution If you upgraded your software after you received your software recovery CD-ROM, using the CD-ROM software images may downgrade your system.

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 WAVE-274/474/574 and WAE-674/7341/7371 devices (64-bit platforms), use the system image named WAAS-4.x.x.x-K9.x86_64.sysimg. For all other devices, use the system image named WAAS-4.x.x.x-K9.sysimg.

These options are available from the software recovery CD-ROM installer menu:

- Option 1: Configure Network—If the .bin image you need to install is located on the network instead of the CD-ROM (which may be the case when an older CD-ROM 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 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 appears only for WAE-7341, WAE-7371, and WAE-674 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.

- Option 9: Exit (Eject and reboot)—This option ejects the CD-ROM and reboots the device. (For all WAE devices except WAE-7341, WAE-7371, and WAE-674 devices, this option is Option 8.)

To reinstall the system software on a WAE appliance using the software recovery CD-ROM, follow these steps:

Step 1 Connect a serial console to the WAE appliance to be upgraded and use the console for the following steps.

Step 2 Insert the WAAS 4.1.x CD-ROM in the CD drive of the WAE device.

Step 3 Reboot the WAE. 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 cdrom
  6. Install flash image from disk
  7. Recreate RAID device (WAE-674/7341/7371 only)
  8. Wipe out disks and install .bin image
  9. Exit (Eject and reboot)
Choice [0]:
```

Option 7 appears only for WAE-7341, WAE-7371, and WAE-674 devices.

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 the CD-ROM.

Step 7 (Optional) If you are working with a WAE-7341, WAE-7371, or WAE-674 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 4.1.x image is installed.

Step 9 Choose the option (8 or 9, depending on your WAE model) to reboot the WAE.

After the WAE reboots, it is running the WAAS 4.1.x 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-4.1.7.10-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-4.1.7.10-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-4.1.1-k9.bin
```

Step 6 Reload the module:

```
NM-WAE-1# reload
```

After the module reboots, it is running the WAAS 4.1.x 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 enable legacy WAFS core or edge services, 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.

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 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 WAVE-274/474/574 and WAE-674/7341/7371 devices (64-bit platforms), use the system image named WAAS-4.x.x.x-K9.x86_64.sysimg. For all other devices, use the system image named WAAS-4.x.x.x-K9.sysimg.

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 device and open a terminal session.

Step 3 Reboot the device by toggling the power on/off switch.

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 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: FastEthernet 0/0
    1: FastEthernet 0/1
    0
Using interface FastEthernet 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 (an email address):

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-4.1.7.10-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-4.1.7.10-K9.sysimg

Attempting download...
Downloaded 10711040 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.
Initializing memory. Please wait.
```

Step 4 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> enable
Console# show version
Wide Area Application Services (WAAS)
Copyright (c) 1999-2008 by Cisco Systems, Inc.
Wide Area Application Services Release 4.1.1
Version: ce507-5.2.0

Compiled 02:34:38 May 8 2008 by (cisco)
Compile Time Options: PP SS



System was restarted on Thu June 22 16:03:51 2008.
```

```
The system has been up for 4 weeks, 1 day, 6 hours, 7 minutes, 23 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 password** command in global configuration mode.

```
WAE# configure
WAE(config)# username admin password 0 password
```

You can specify that the password be either cleartext or encrypted.

Note Do not set the user ID (uid).

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 automaticaly 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 GUI navigation pane, go to **My WAN > Manage Devices**.

b. Click the **Edit** icon next to the device that you want to deactivate.

c. From the navigation pane, choose *Device Name* **> Activation**. The Device Activation window appears.

d. 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, 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 15-21.

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 WAAS GUI, choose **My WAN > Manage Devices**.

b. Click the **Edit** icon next to the device that you want to deactivate. The Device Dashboard window appears.

c. In the navigation pane, choose *Device Name* **> Activation**.

d. Uncheck the **Activate** check box. The window refreshes, displaying a check box for marking the device as replaceable.

e. Check the **Replaceable** check box, and click **Submit**.

Note This check box appears in the GUI only when the device is inactive.

Step 2 Configure a system device recovery key as follows:

a. From the WAAS Central Manager GUI navigation pane, choose **Configure > 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 3 Configure the basic network settings for the new device.

Step 4 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 5 Activate the device:

a. From the WAAS GUI, choose **My WAN > Manage Devices**.

b. Click the **Edit** icon next to the WAAS device that you want to activate. The Device Dashboard window appears.

c. In the navigation pane, choose *Device Name* **> Activation**. The WAAS device status should be Online.

d. Check the **Activate** check box, and click **Submit**.

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 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.

Table 15-3 shows the drive-type compatibility for the WAE-612. All drives must be the same type.

Table 15-3 WAE-612 Drive-Type Compatibility Matrix

Drive Types	SAS[1]	SATA2[2]
SAS	Ok	No
SATA2	No	Ok

1. Serial Attached SCSI
2. Serial Advanced Technology Attachment 2

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.

Note Before you configure a standby Central Manager, you must manually install your print drivers if you are using WAAS print services. Print drivers are not automatically replicated from the primary Central Manager database to the standby device.

This section contains the following topics:

- Converting a WAE to a Standby Central Manager, page 15-25
- Converting a Primary Central Manager to a Standby Central Manager, page 15-26
- Converting a Standby Central Manager to a Primary Central Manager, page 15-26
- Switching Both Central Manager Roles, page 15-27

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 and Application Accelerator functionality.
- Accelerator only—Includes Application Accelerator functionality only. If you want to change an Application Accelerator 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 15-1.

You can use the **show version** EXEC command to check if the WAE is running an Accelerator only image; if it is, the output displays "(WAAS-ACCELERATOR-K9)." Also, the **show running-config** EXEC command displays similar output if the WAE is running an Accelerator only image.

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 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 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 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
```

If you are switching Central Manager roles, see the next section, "Switching Both Central Manager Roles" section on page 15-27.

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.

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 **My WAN > Manage Devices**, choose a device, 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.

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device (or device group) for which you want to configure the disk error handling method.

Step 3 In the navigation pane, choose **Configure > Storage > Disk Error Handling**.

The Disk Error Handling Settings window appears.

Step 4 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 5 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 only supported on WAE-674-4G and 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device (or device group) for which you want to enable extended object cache.

Step 3 In the navigation pane, choose **Configure > Storage > Extended Object Cache**.

The Extended Disk Space Settings window appears.

Step 4 Check the **Enable** check box to enable extended object cache. See Table 15-4 and Table 15-5 for sizing details.

Note If extended object cache is enabled on WAE-674-4G or WAE-674-8G models with a vbspace that is greater than 30-GB, the device goes into override mode. In override mode, a force DG settings button will be present on the device page (or device group page).

If you have a virtual blade enabled using a vbspace 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 vbspace filesystem will be reduced to 30 GB.

Step 5 Click **Submit** to save the settings.

The disk cache size on the WAE-674 platform varies depending upon the features enabled.

Table 15-4 contains the disk cache sizes for the WAE-674-4G platform.

Table 15-4 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	120GB	120GB	120GB	120GB
CIFS Object Cache	95GB	95GB	340GB	300GB
Virtual Blade	--	120GB	--	30GB

Table 15-5 contains the disk cache sizes for the WAE-674-8G platform.

Table 15-5 WAE-674-8G 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	320GB	150GB	150GB	150GB
CIFS Object Cache	95GB	95GB	310GB	275GB
Virtual Blade	--	200GB	--	30GB

Activating All Inactive WAAS Devices

To activate all inactivated WAAS devices in your network, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**. The Devices listing window appears.

Step 2 Click the **Activate all inactive WAEs** icon in the taskbar. The Activate All Inactive WAEs window appears. (See Figure 15-2.)

Figure 15-2 Activating Inactive Devices

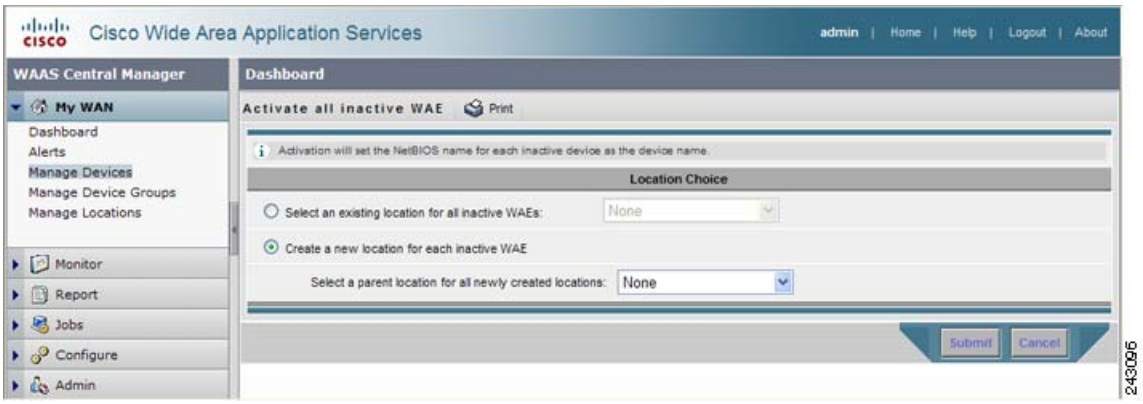

Step 3 Choose an existing location for all inactivated WAAS devices by clicking the **Select an existing location for all inactive WAEs** radio button, and then choose a location from the drop-down list.

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 GUI, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the name of the device that you want to reboot. The Device Dashboard window appears.

Step 3 In the taskbar, click the **Reload WAE** icon. You are prompted to confirm your decision.

Step 4 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, 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 GUI, choose **My WAN > Manage Device Groups**.

Step 2 Click the **Edit** icon next to the name of the device group that you want to reboot. The Modifying Device Group window appears.

Step 3 In the taskbar, click the **Reboot All Devices in Device Group** icon. You are prompted to confirm your decision.

Step 4 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 16

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.

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, page 16-2
- Troubleshooting Devices Using Alerts, page 16-6
- Viewing Device Information, page 16-7
- Customizing a Dashboard or Report, page 16-10
- Chart Descriptions, page 16-14
- Using Predefined Reports to Monitor WAAS, page 16-35
- Managing Reports, page 16-48
- Configuring Flow Monitoring, page 16-53
- Configuring System Logging, page 16-55
- Configuring Transaction Logging, page 16-58
- Viewing the System Message Log, page 16-61
- Viewing the Audit Trail Log, page 16-62
- Viewing the Device Log, page 16-63
- Enabling the Kernel Debugger, page 16-63
- Troubleshooting Using Diagnostic Tests, page 16-64
- Using the show and clear Commands from the WAAS Central Manager GUI, page 16-65

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:

- Monitoring Graphs and Charts, page 16-2
- Alarm Panel, page 16-3
- Device Alarms, page 16-5

Figure 16-1 shows the System Dashboard window.

Figure 16-1 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 > 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 four tabs that show several graphical displays about the application traffic processed by your WAAS system:

- The Traffic tab contains the following graphical displays:
  - Traffic Summary chart—Displays the applications with the highest percentage of traffic in the WAAS network for the last hour.
  - Original Traffic over Time graph—Displays the amount of original and pass-through traffic over the last hour.

  - Traffic Volume and Reduction graph—Displays the original and optimized traffic volume and percentage of traffic reduction over the last hour.
  - Optimized Traffic over Time graph—Displays the amount of optimized and pass-through traffic over the last hour. (This graph is minimized.)
- The Optimization tab is the default tab and contains the following graphical displays:
  - 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.
  - Bandwidth Optimization chart—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 chart—Displays the volume of original and optimized traffic, and the percentage of traffic reduction, for the last hour.
- The Acceleration tab contains the following graphical displays:
  - HTTP: Estimated Time Savings graph—Displays an estimate of the amount of response time saved for HTTP traffic for the last hour due to the HTTP accelerator.
  - MAPI: Estimated Time Savings graph—Displays an estimate of the amount of response time saved for MAPI traffic for the last hour due to the MAPI accelerator.
  - NFS: Estimated Time Savings graph—Displays an estimate of the amount of response time saved for NFS traffic for the last hour due to the NFS accelerator.
- The Platform tab contains the following graphical display:
  - Managed Devices Information—Displays the total number of WAAS devices in your network, showing the number that are online, offline, inactive, and pending. This display also shows the number of alarms of different severities and the software versions that are deployed in your network.

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 16-10. Individual charts are described in more detail in the "Chart Descriptions" section on page 16-14.

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 9-5. 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 in the System Dashboard window provides a near real-time view of incoming alarms. The panel refreshes every 2 minutes to reflect updates to the system alarm database.

The alarm panel contains two tabs: Active Alarms and Acknowledged Alarms. The Active Alarms tab displays a dynamic view of all incoming alarms. You may remove an alarm from the active display by acknowledging the alarm. Acknowledged alarms are moved to the Acknowledged Alarms view. You may choose to unacknowledged an alarm and return it to the Active view at any time.

Only Active alarms can be acknowledged in the alarm panel. Pending, Offline, and Inactive alarms cannot be acknowledged in the alarm panel.

For either view, 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.

When you roll your mouse over an item under the Alarm Information column, a contextual popup menu appears. The popup menu provides links to the troubleshooting and monitoring windows in the WAAS Central Manager GUI. For more information on these links, see the "Troubleshooting Devices Using Alerts" section on page 16-6.

Figure 16-2 shows the alarm panel in the System Dashboard window.

Figure 16-2 System Dashboard Window Alarm Panel

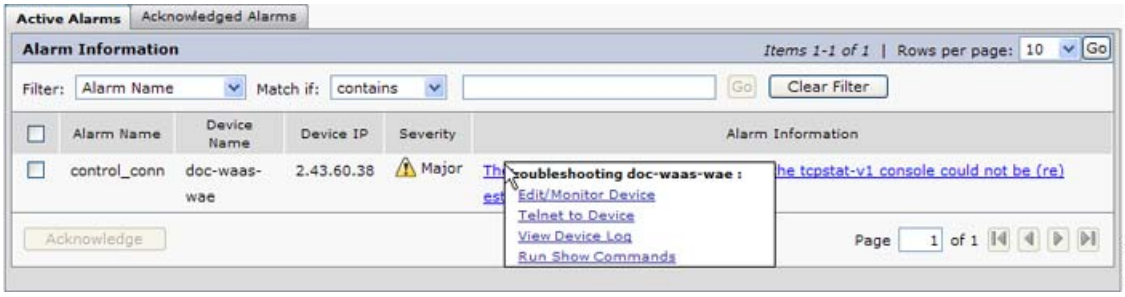

To acknowledge an active alarm and move it from Active Alarms to the separate Acknowledged Alarms section, follow these steps:

Step 1 From the System Dashboard window alarm panel, check the check box next to the name of the alarm that you want to acknowledge.

Step 2 Click the **Acknowledge** button.

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 Active 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.

The alarm will be moved to the Acknowledged Alarms tab.

To filter and sort alarms displayed in the System Dashboard window alarm panel, follow these steps:

Step 1 From the Filter drop-down list, choose one of the following filtering options:

- **Alarm Name**
- **Device Name**

- **Device IP**
- **Severity**
- **Alarm Information**

Step 2 From the Match if drop-down list, choose one of the following match conditions:

- **contains**
- **starts with**
- **ends with**
- **is exactly**
- **not exactly**
- **not contain**
- **clear**

Step 3 Enter a match string in the text entry field. This field accepts any alphanumeric text, including special characters.

Step 4 Click **Go**.

Step 5 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 6 To clear the filter, click **Clear**.

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 16-1 describes the various device alarms that can appear.

Table 16-1 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.
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 GUI navigation pane, invoke the Troubleshooting Devices window in one of these ways:

- Choose **My WAN > Alerts** to view alarms on all devices.
- Choose **My WAN > Manage 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 16-3.)

Figure 16-3 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 16-2 describes the tools available for device alarms.

Table 16-2 Troubleshooting Tools for Device Alarms

Item	Navigation	Description
Update Software	Choose device, **Jobs > 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.

Table 16-2 Troubleshooting Tools for Device Alarms (continued)

Item	Navigation	Description
Telnet to Device	Opens a Telnet window	Initiates a Telnet session using the device IP address.
View Device Log	Choose device, **Admin > Logs**	Displays system message logs filtered for this device.
Run Show Commands	Choose device, **Troubleshoot > CLI Commands > Show Commands**	Displays device **show** command tool. For more information, see the "Using the show and clear Commands from the WAAS Central Manager GUI" section on page 16-65.

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 **My WAN > Manage Devices** in the WAAS Central Manager GUI.

Figure 16-4 shows an example of the Devices window.

Figure 16-4 Devices Window

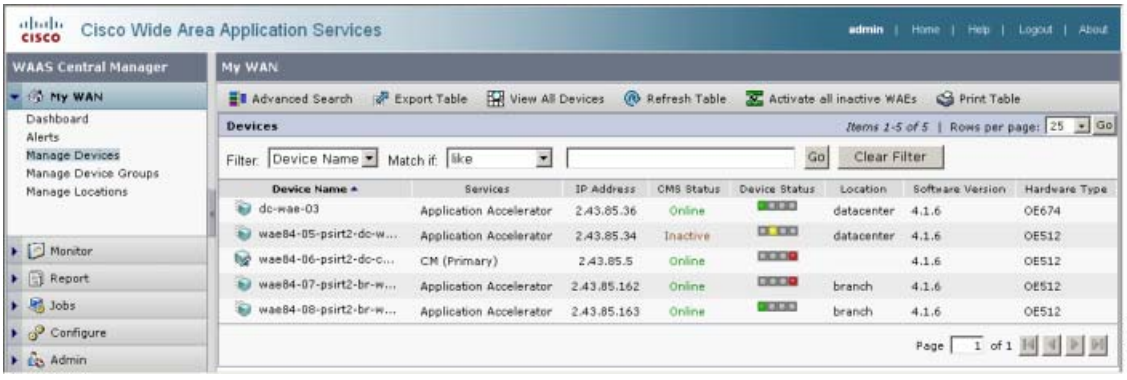

This window displays the following information about each device:

- Services enabled on the device. See Table 16-3 for a description of these services.

- IP address of the device.
- CMS Status (Online, Offline, Pending, or Inactive). For more information about the status, see the "Device Alarms" section on page 16-5.
- 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 16-6.
- 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 16-36).
- Software version installed and running on the device.
- Hardware type of the device. If you see a type such as OE574, the numbers refer to the model number, WAVE-574 in this example. NME-WAE refers to a NME-WAE module and SM-WAE refers to a SM-SRE module.

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 16-3 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 15-26.
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 15-26.
Application Accelerator	The device has been enabled as an application accelerator.
Replication Accelerator	The device has been enabled as a replication accelerator. (Supported only on 4.0.19 or later 4.0.x devices.)
Edge	The device has been enabled with legacy WAFS Edge services so it can accelerate data stored on a remote file server. For information on enabling WAFS, see Chapter 11, "Configuring Wide Area File Services."

Table 16-3 Service Descriptions (continued)

Service	Description
Core	The device has been enabled with legacy WAFS Core services so it can accelerate data stored on a remote file server. For information on enabling WAFS, see Chapter 11, "Configuring Wide Area File Services."
Print	The device has been enabled with legacy print services so it can act as a print server to branch office clients. For information on setting up a print server, see Chapter 13, "Configuring and Managing WAAS Legacy Print Services."

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 16-5.)

To access the Device Dashboard window, go to **My WAN > Manage Devices** and click the **Edit** icon next to the device that you want to view.

Figure 16-5 Device Dashboard Window

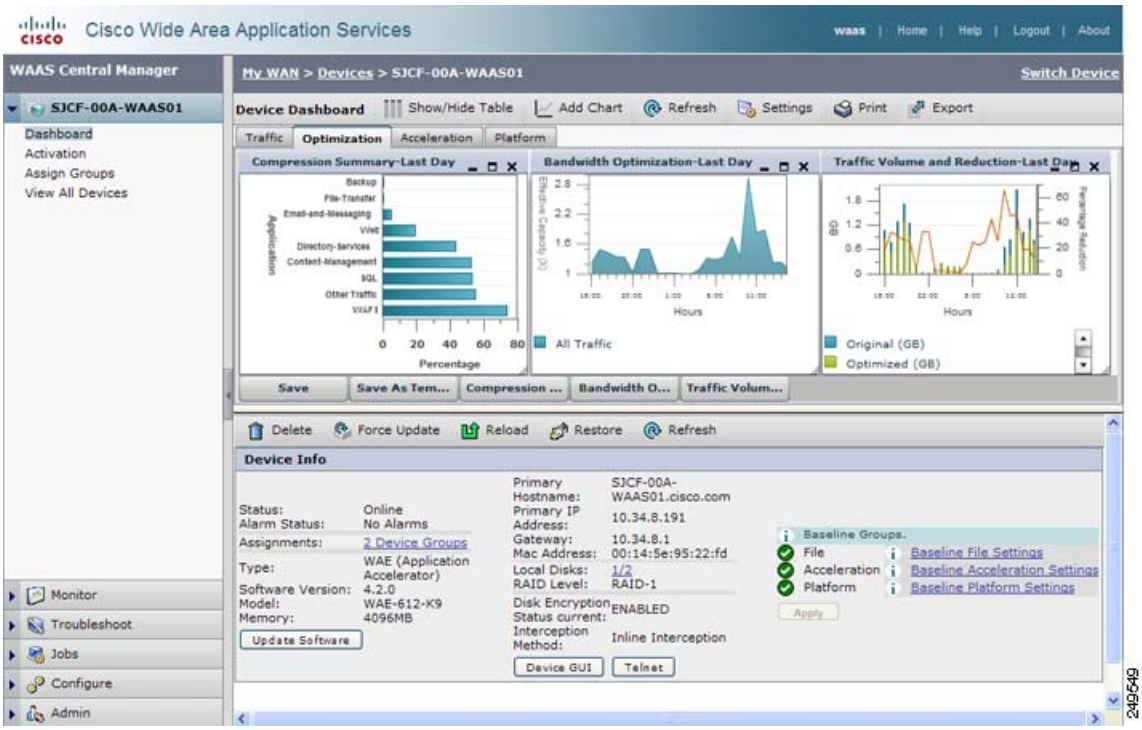

From the Device Dashboard window, you can perform the following tasks:

- 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 chart panel in the top of the window. For more information, see the "Customizing a Dashboard or Report" section on page 16-10. Individual charts are described in more detail in the "Chart Descriptions" section on page 16-14.

- 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.
- 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 16-10.
- Click **Update Software** to update the software on the device. For more information, see Chapter 15, "Maintaining Your WAAS System."
- Click **Device GUI** to open the WAE Device Manager. For more information on managing a device using this GUI, see Chapter 10, "Using the WAE Device Manager GUI."
- Click **Telnet** to establish a Telnet session into the device and issue CLI commands.
- Assign and unassign the device to baseline groups. For more information, see Chapter 3, "Using Device Groups and Device Locations."

Viewing and Unlocking Device Users

To view the users that are defined on a device, go to **My WAN > Manage Devices** and click the **Edit** icon next to the device that you want to view. Then from the navigation pane, 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. The dashboards and reports work the same way. For more information about creating custom reports, see the "Managing Reports" section on page 16-48.

If the charts panel (see Figure 16-6) in the upper part of the window has tabs, you can click other tabs to display other groups of charts.

You can use the minimize, maximize, and close buttons on the top of individual chart panes to minimize, maximize, and close the chart. When a chart is maximized, the middle button changes to a restore button that restores the chart to its smaller size. You can move a chart pane by clicking and dragging its title bar. You can resize a chart by clicking and dragging on the resize control in the lower right corner of the chart pane.

Figure 16-6 Charts Panel

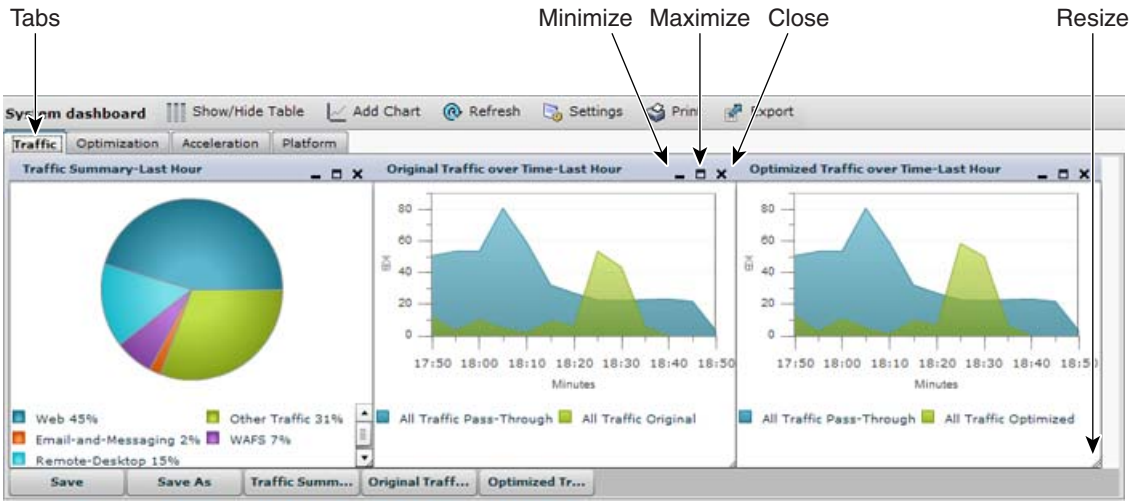

Icons across the top of the dashboard allow you to do the following:

- Show/Hide Table—Shows or hides the table panel in the lower part of the window. For the system dashboard, this is the Alarms table, and for the device dashboard, it is the Device Information table. Dashboard tables are fixed and cannot be changed. You can include different tables in custom reports.
- Add Chart—Adds a graphical chart to the charts panel. You can display a maximum of six charts. For details on adding a chart, see the "Adding a Chart" section on page 16-12. Individual charts are described in more detail in the "Chart Descriptions" section on page 16-14.
- Refresh—Refreshes the charts with new information.
- Settings—Sets the time frame and included applications for the data presented in each chart. For details on these settings, see the "Configuring Chart Settings" section on page 16-12.
- Print—Prints 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.

Buttons below the charts panel allow you to do the following:

- Save—Saves the dashboard or report with its current settings. The next time you view it, it is displayed with these settings.
- Save As Template—Saves the report with its current settings under a new name. A pop-up window lets you enter a report name and optional notes about the report. Only the following characters are allowed to be entered: numbers, letters, spaces, periods, hyphens, and underscores. The report will be available in the **Report > Manage Reports** window.
- Chart Names—Displays or hides the named chart.

Adding a Chart

To add a chart to a dashboard or report, follow these steps:

Step 1 From the dashboard or report chart panel, click the **Add Chart** icon in the taskbar. The Add Chart window is displayed, as shown in Figure 16-7.

Figure 16-7 Adding a Chart

Step 2 Expand any of the chart categories by clicking on the plus sign next to the category.

Step 3 Check the box next to each chart that you want to display. Individual charts are described in more detail in the "Chart Descriptions" section on page 16-14.

Remember that a report can contain a maximum of six charts.

Step 4 Click **Add**.

If you want to delete a chart 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 From the dashboard or report chart panel, click the **Settings** icon in the taskbar. The Settings window is displayed, as shown in Figure 16-8.

Figure 16-8 Chart Settings

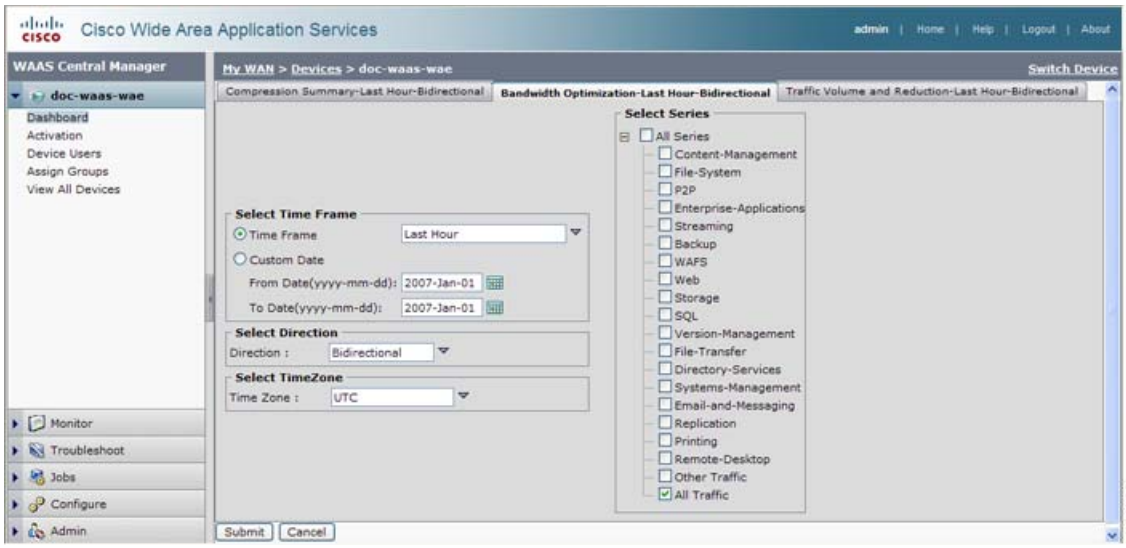

Step 2 Click on the tab corresponding to the chart that you want to change.

Step 3 In the Select Time Frame area, choose a time frame for the chart. Choose one of these options by clicking on the appropriate radio button:

- **Time Frame**—Choose one of 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 9-17.
  - **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 Date**—Enter starting and ending dates in the From Date and To Date fields. Click the calendar icon to choose dates from a popup calendar.

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. This behavior may give the chart the appearance of 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 is normal.

Step 4 From the 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.

The Select Direction area appears only when you are configuring certain device level charts.

Step 5 From the Time Zone drop-down list, choose one of the following options:

- **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).
- **WAE Local Time**—Sets the time zone of the report to the time zone of the WAE device. This option appears only when you are configuring a device level chart.

Step 6 In the Select Series area, click the plus sign to expand the All Series list. Place a check mark next to the applications whose statistics you want to include in the chart data. This optional area applies only to specific chart types. To include all applications, choose **All Traffic** (default).

Note When more than three applications are selected, area comparison charts are converted to line charts for easier readability. This is only applicable to charts with user-selectable applications.

Step 7 Click **Submit**.

Chart Descriptions

This section describes the charts that you can choose to include in a dashboard or report. The following categories of charts are available:

- Traffic Analysis Charts, page 16-14
- Optimization Charts, page 16-15
- Acceleration Charts, page 16-18
- Platform Charts, page 16-34

All charts are plotted using the Central Manager local time zone, unless the chart settings were customized to use a different time zone.

Traffic Analysis Charts

This section describes these charts:

- Traffic Summary, page 16-14
- Original Traffic Over Time, page 16-15

Traffic Summary

The Traffic Summary chart (see Figure 16-9) displays the top six applications with the highest percent of traffic. Each section in the pie chart represents an application as a percent 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 seventh category named Other Traffic.

Figure 16-9 Traffic Summary Chart

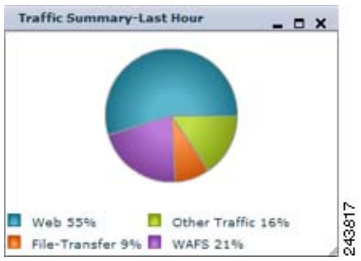

Formula:

(App Traffic/ Total Traffic) * 100

App Traffic is the Original traffic (Original Excluding Pass-Through + Pass-Through) flowing for an application.

Original Traffic Over Time

The Original Traffic over Time chart (see Figure 16-10) displays a graph of the amount of original and pass-through traffic. You can choose which applications to include; the default is all traffic. The chart in the display is an area chart by default. If you select more than three applications a line chart is used for easier readability.

Figure 16-10 Original Traffic Over Time Chart

Optimization Charts

This section describes these charts:

- Compression Summary, page 16-16
- Compression Over Time, page 16-16
- Compression by Application Over Time, page 16-16
- Optimized Traffic Over Time, page 16-17
- Traffic Volume and Reduction, page 16-17
- Bandwidth Optimization, page 16-18

Compression Summary

The Compression Summary chart (see Figure 16-11) displays a bar chart of the percentage of traffic reduction (excluding pass-through traffic) for the top ten applications with the highest amount of traffic.

Figure 16-11 Compression Summary Chart

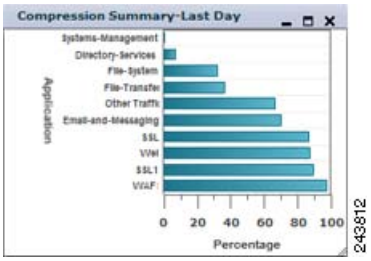

Formula:

% Reduction Excluding Pass-Through = (Original Excluding Pass-Through – Optimized) / (Original Excluding Pass-Through)

Compression Over Time

The Compression over Time chart (see Figure 16-12) 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.

Figure 16-12 Compression Over Time Chart

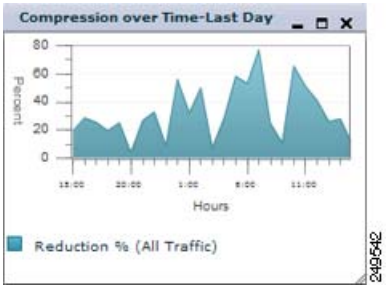

Formula:

% Reduction = (Original Excluding Pass-Through – Optimized) / (Original Excluding Pass-Through)

Compression by Application Over Time

The Compression by Application over Time chart (see Figure 16-13) displays a graph of the percent of total traffic that was reduced on the WAE device using the WAAS optimization techniques. This chart excludes pass-through traffic in the results. You can choose which applications to include; the default is all traffic, which shows the same information as the Compression Over Time chart. The chart in the display is an area chart by default. If you select more than three applications a line chart is used for easier readability.

Figure 16-13 Compression By Application Over Time Chart

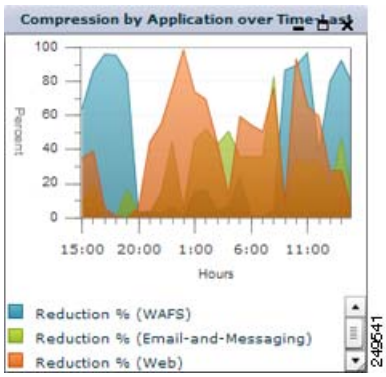

Formula:

% Reduction = (Original Excluding Pass-Through – Optimized) / (Original Excluding Pass-Through)

Optimized Traffic Over Time

The Optimized Traffic over Time chart (see Figure 16-14) displays a graph of the amount of optimized and pass-through traffic processed on the WAE device. The traffic units (bytes, KB, MB, or GB) at the left side depend upon the range. The chart in the display is an area chart by default. If you select more than three applications a line chart is used for easier readability.

Figure 16-14 Optimized Traffic Over Time Chart

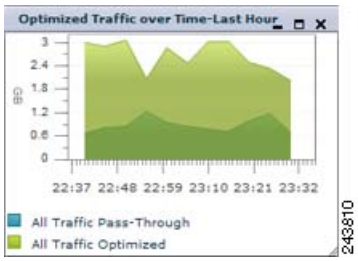

Traffic Volume and Reduction

The Traffic Volume and Reduction chart (see Figure 16-15) displays the amount of original and optimized traffic in a bar chart, along with the percentage of traffic reduction (excluding pass-through traffic) as a line. The traffic units (bytes, KB, MB, or GB) at the left side depend upon the range. The percentage of traffic reduction units are shown at the right side of the chart.

Figure 16-15 Traffic Volume and Reduction Chart

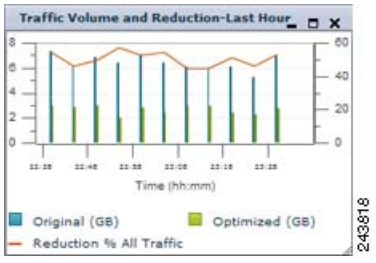

Formula:

% Reduction Excluding Pass-Through = (Original Excluding Pass-Through – Optimized) / (Original Excluding Pass-Through)

Bandwidth Optimization

The Bandwidth Optimization chart (see Figure 16-16) 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. The chart in the display is an area chart by default. If you select more than three applications a line chart is used for easier readability.

Figure 16-16 Bandwidth Optimization Chart

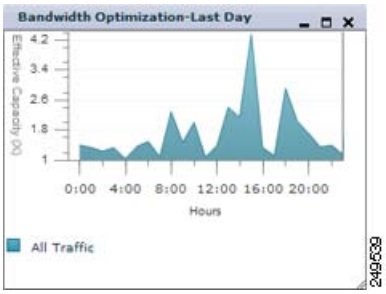

Formula:

Effective WAN Capacity = 1 / (1-% Reduction Excluding Pass-Through)

% Reduction Excluding Pass-Through = (Original Excluding Pass-Through – Optimized) / (Original Excluding Pass-Through)

Acceleration Charts

This section describes these charts:

- HTTP, page 16-18
- CIFS, page 16-21
- MAPI, page 16-25
- NFS, page 16-28
- Video, page 16-31
- SSL, page 16-32

HTTP

This section describes these charts:

- HTTP: Estimated Time Savings, page 16-19
- HTTP: Connection Details, page 16-19
- HTTP: Bandwidth Optimization, page 16-19
- HTTP: Response Time Savings, page 16-20

- HTTP: Optimization Count, page 16-20
- HTTP: Optimization Techniques, page 16-21

HTTP: Estimated Time Savings

The HTTP Estimated Time Savings chart (see Figure 16-17) displays a graph of the estimated percentage of response time saved by the HTTP accelerator due to fast connection reuse and metadata caching.

Figure 16-17 HTTP Estimated Time Savings Chart

HTTP: Connection Details

The HTTP Connection Details chart (see Figure 16-18) displays the HTTP session connection statistics, showing the total number of connections handled with the number of unaccelerated (bypassed) connections. The chart in the display is an area chart; the total number of connections handled is indicated by the color blue and is shown behind the number of bypassed connections, which is indicated by the color green.

Figure 16-18 HTTP Connection Details Chart

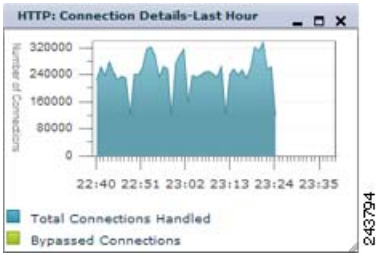

HTTP: Bandwidth Optimization

The HTTP Bandwidth Optimization chart (see Figure 16-19) displays the effective bandwidth capacity of the WAN link as a result of HTTP acceleration, as a multiplier of its base capacity. The chart in the display is an area chart; data for all traffic is indicated by the color blue and is shown behind web (HTTP) traffic, which is indicated by the color green.

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.

Figure 16-19 HTTP Bandwidth Optimization Chart

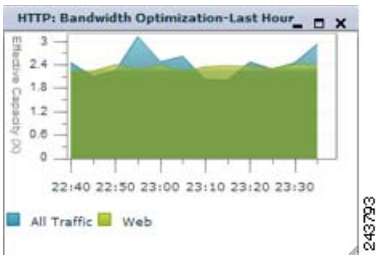

HTTP: Response Time Savings

The HTTP Response Time Savings chart (see Figure 16-20) 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 chart in the display is a stacked area chart. The time units (milliseconds, seconds, or minutes) at the left side depend upon the range.

Figure 16-20 HTTP Response Time Savings Chart

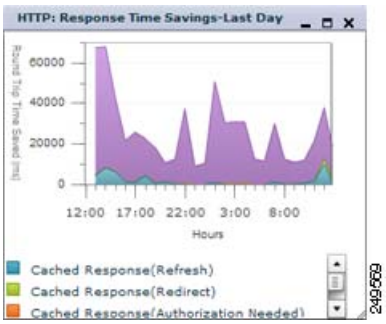

HTTP: Optimization Count

The HTTP Optimization Count chart (see Figure 16-21) 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.

Figure 16-21 HTTP Optimization Count Chart

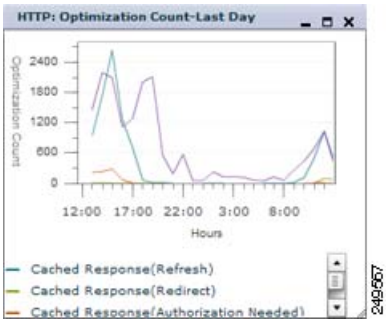

HTTP: Optimization Techniques

The HTTP Optimization Techniques pie chart (see Figure 16-22) 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.

Figure 16-22 HTTP Optimization Techniques Chart

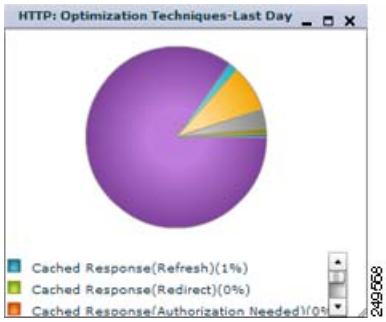

CIFS

This section describes these charts:

- CIFS: Request Optimization, page 16-21
- CIFS: Cached Objects, page 16-22
- CIFS: Cache Utilization, page 16-22
- CIFS: Connection Statistics, page 16-23
- CIFS: File Optimization, page 16-23
- CIFS: Client Average Throughput, page 16-23
- CIFS: Connected CIFS Core Count, page 16-24
- CIFS: CIFS Edge-CIFS Core Traffic, page 16-24
- CIFS: Connected CIFS Edge Count, page 16-25
- CIFS: CIFS Core Traffic, page 16-25

All CIFS charts are available only at the device level and the specific charts available depend on the mode of the device: transparent CIFS accelerator mode, WAFS legacy mode as an Edge device, or WAFS legacy mode as a Core device.

CIFS: Request Optimization

The CIFS Request Optimization chart (see Figure 16-23) displays the percentage of requests that are served locally from the CIFS cache. This chart is available only for a device that is operating in transparent CIFS accelerator mode or WAFS legacy mode as an Edge device.

Figure 16-23 CIFS Request Optimization Chart

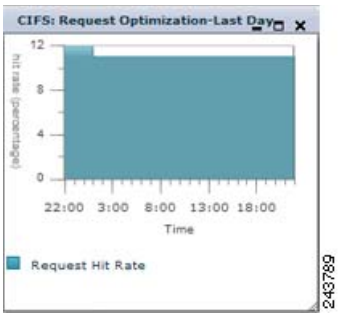

CIFS: Cached Objects

The CIFS Cached Objects chart (see Figure 16-24) displays the number of objects in the CIFS cache. This chart is available only for a device that is operating in transparent CIFS accelerator mode or WAFS legacy mode as an Edge device.

Figure 16-24 CIFS Cached Objects Chart

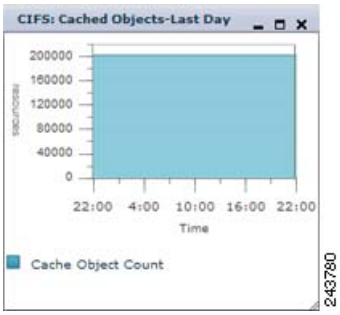

CIFS: Cache Utilization

The CIFS Cache Utilization chart (see Figure 16-25) displays the utilization percentage of the CIFS cache. This chart is available only for a device that is operating in transparent CIFS accelerator mode or WAFS legacy mode as an Edge device.

Figure 16-25 CIFS Cache Utilization Chart

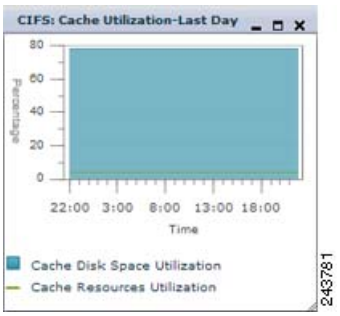

CIFS: Connection Statistics

The CIFS Connection Statistics chart (see Figure 16-26) displays the number of CIFS accelerated sessions. This chart is available only for a device that is operating in transparent CIFS accelerator mode or WAFS legacy mode as an Edge device.

Figure 16-26 CIFS Connection Statistics Chart

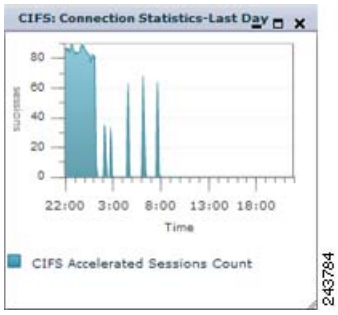

CIFS: File Optimization

The CIFS File Optimization chart (see Figure 16-27) displays the number of open CIFS files. This chart is available only for a device that is operating in transparent CIFS accelerator mode or WAFS legacy mode as an Edge device.

Figure 16-27 CIFS File Optimization Chart

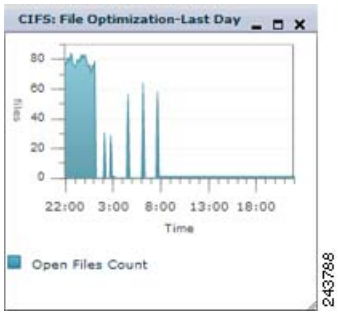

CIFS: Client Average Throughput

The CIFS Client Average Throughput chart (see Figure 16-28) displays the average throughput (in KB/second) between the WAAS device and its clients. This chart is available only for a device that is operating in transparent CIFS accelerator mode or WAFS legacy mode as an Edge device.

Figure 16-28 CIFS Client Average Throughput Chart

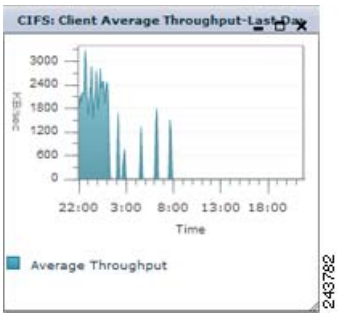

CIFS: Connected CIFS Core Count

The CIFS Connected CIFS Core Count chart (see Figure 16-29) displays the number of connected CIFS Core devices. This chart is available only for a device that is operating in WAFS legacy mode as an Edge device.

Figure 16-29 CIFS Connected CIFS Core Count Chart

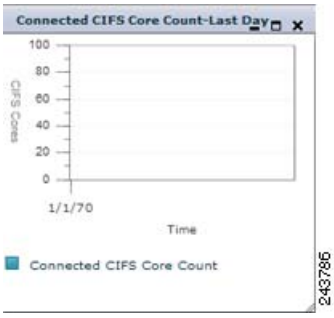

CIFS: CIFS Edge-CIFS Core Traffic

The CIFS Edge-CIFS Core Traffic chart (see Figure 16-30) displays the amount of traffic between the Edge device and the CIFS core devices that are connected to it. This chart is available only for a device that is operating in WAFS legacy mode as an Edge device.

Figure 16-30 CIFS Edge-CIFS Core Traffic Chart

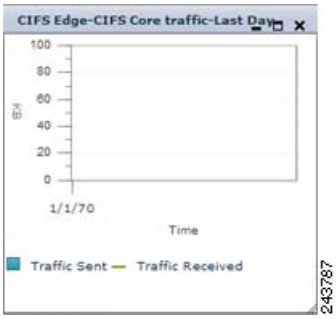

CIFS: Connected CIFS Edge Count

The CIFS Connected CIFS Edge Count chart (see Figure 16-31) displays the number of connected CIFS Edge devices. This chart is available only for a device that is operating in WAFS legacy mode as a Core device.

Figure 16-31 CIFS Connected CIFS Edge Count Chart

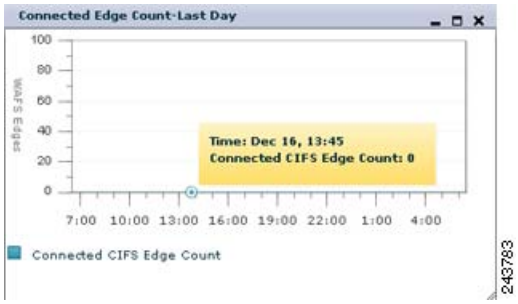

CIFS: CIFS Core Traffic

The CIFS Core Traffic chart (see Figure 16-32) displays the amount of traffic between the Core device and the CIFS Edge devices that are connected to it. This chart is available only for a device that is operating in WAFS legacy mode as a Core device.

Figure 16-32 CIFS Core Traffic Chart

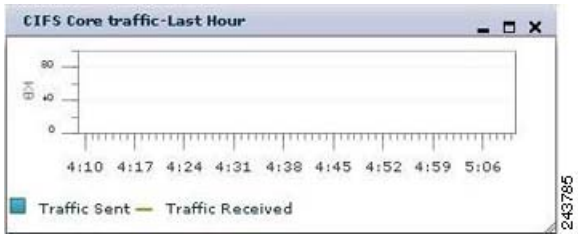

MAPI

This section describes these charts:

- MAPI: Request Optimization, page 16-26
- MAPI: Response Time Optimization, page 16-26
- MAPI: Versions Detected, page 16-26
- MAPI: Estimated Time Savings, page 16-26
- MAPI: Connection Details, page 16-27
- MAPI: Acceleration Bypass Reason, page 16-27
- MAPI: Bandwidth Optimization, page 16-28

MAPI: Request Optimization

The MAPI Request Optimization chart (see Figure 16-33) displays the percentage of local and remote MAPI command responses. A local response is one that is sent to the client without waiting for a response from the peer WAE. A remote response comes from the remote server.

Figure 16-33 MAPI Request Optimization Chart

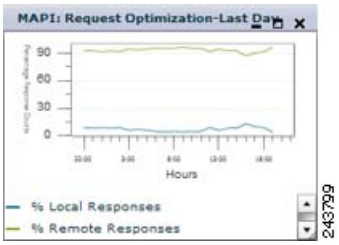

MAPI: Response Time Optimization

The MAPI Response Time Optimization chart (see Figure 16-34) 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.

Figure 16-34 MAPI Response Time Optimization Chart

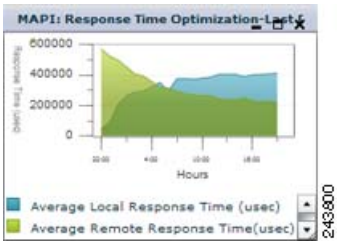

MAPI: Versions Detected

The MAPI Versions Detected pie chart (see Figure 16-35) displays the number of connections detected from different versions (2000, 2003, 2007) of the Microsoft Outlook client.

Figure 16-35 MAPI Versions Detected Chart

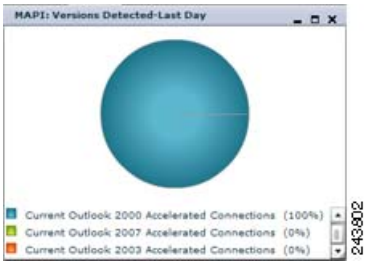

MAPI: Estimated Time Savings

The MAPI Estimated Time Savings chart (see Figure 16-36) displays a graph of the estimated percentage of response time saved by the MAPI accelerator.

Figure 16-36 MAPI Estimated Time Savings Chart

MAPI: Connection Details

The MAPI Connection Details chart (see Figure 16-37) displays the MAPI session connection statistics, showing the total number of connections handled with the number of unaccelerated (bypassed) connections. The chart in the display is an area chart; the total number of connections handled is indicated by the color blue and is shown behind the number of bypassed connections, which is indicated by the color green.

Figure 16-37 MAPI Connection Details Chart

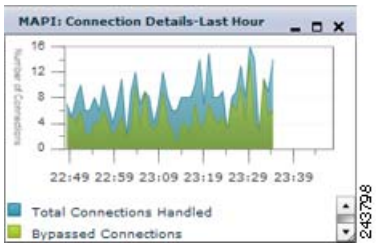

MAPI: Acceleration Bypass Reason

The MAPI Acceleration Bypass Reason pie chart (see Figure 16-38) displays the reasons that MAPI traffic is not accelerated: 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.

Figure 16-38 MAPI Acceleration Bypass Reason Chart

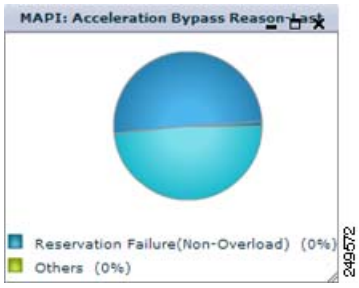

MAPI: Bandwidth Optimization

The MAPI Bandwidth Optimization chart (see Figure 16-39) displays the effective bandwidth capacity of the WAN link as a result of MAPI acceleration, as a multiplier of its base capacity. The chart in the display is an area chart; data for all traffic is indicated by the color blue and is shown behind e-mail and messaging (MAPI) traffic, which is indicated by the color green.

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.

Figure 16-39 MAPI Bandwidth Optimization Chart

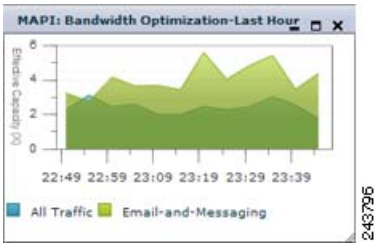

NFS

This section describes these charts:

- NFS: Request Optimization, page 16-28
- NFS: Response Time Optimization, page 16-29
- NFS: Versions Detected, page 16-29
- NFS: Estimated Time Savings, page 16-29
- NFS: Connection Details, page 16-30
- NFS: Acceleration Bypass Reason, page 16-30
- NFS: Bandwidth Optimization, page 16-30

NFS: Request Optimization

The NFS Request Optimization chart (see Figure 16-40) displays the percentage of local and remote NFS command responses. A local response is one that is sent to the client without waiting for a response from the peer WAE. A remote response comes from the remote server.

Figure 16-40 NFS Request Optimization Chart

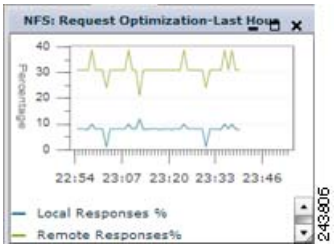

NFS: Response Time Optimization

The NFS Response Time Optimization chart (see Figure 16-41) 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.

Figure 16-41 NFS Response Time Optimization Chart

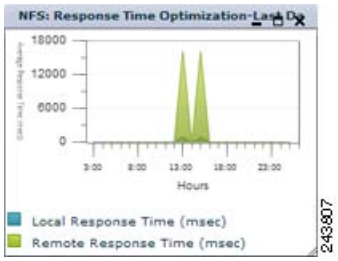

NFS: Versions Detected

The NFS Versions Detected pie chart (see Figure 16-42) 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.

Figure 16-42 NFS Versions Detected Chart

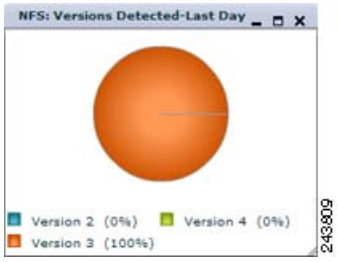

NFS: Estimated Time Savings

The NFS Estimated Time Savings chart (see Figure 16-43) displays a graph of the estimated percentage of response time saved by the NFS accelerator.

Figure 16-43 NFS Estimated Time Savings Chart

NFS: Connection Details

The NFS Connection Details chart (see Figure 16-44) displays the NFS session connection statistics, showing the total number of connections handled with the number of unaccelerated (bypassed) connections. The chart in the display is an area chart; the total number of connections handled is indicated by the color blue and is shown behind the number of bypassed connections, which is indicated by the color green.

Figure 16-44 NFS Connection Details Chart

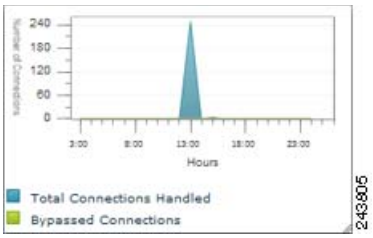

NFS: Acceleration Bypass Reason

The NFS Acceleration Bypass Reason pie chart (see Figure 16-45) displays the reasons that NFS traffic is not accelerated: unknown authentication flavor or unknown NFS version.

Figure 16-45 NFS Acceleration Bypass Reason Chart

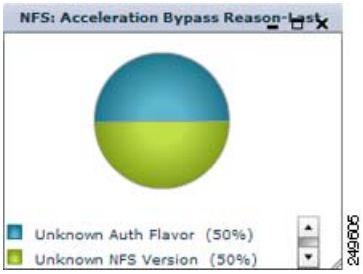

NFS: Bandwidth Optimization

The NFS Bandwidth Optimization chart (see Figure 16-46) displays the effective bandwidth capacity of the WAN link as a result of NFS acceleration, as a multiplier of its base capacity. The chart in the display is an area chart; data for all traffic is indicated by the color blue and is shown behind file system (NFS) traffic, which is indicated by the color green.

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.

Figure 16-46 NFS Bandwidth Optimization Chart

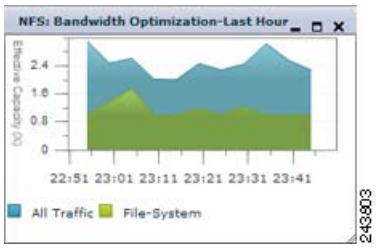

Video

This section describes these charts:

- Video: Stream Optimization, page 16-31
- Video: Connection Details, page 16-31
- Video: Acceleration Bypass Reason, page 16-32
- Video: Bandwidth Optimization, page 16-32

Video: Stream Optimization

The Video Stream Optimization chart (see Figure 16-47) compares the amounts of incoming and outgoing traffic. The chart in the display is an area chart; the total incoming bytes is indicated by the color blue and is shown behind the total outgoing bytes, which is indicated by the color green. The traffic units (bytes, KB, MB, or GB) at the left side depend upon the range.

Figure 16-47 Video Stream Optimization Chart

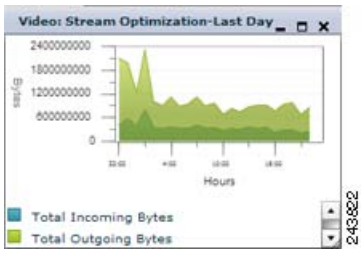

Video: Connection Details

The Video Connection Details chart (see Figure 16-48) displays the video session connection statistics, showing the total number of connections handled with the number of unaccelerated (bypassed) connections. The chart in the display is an area chart; the total number of connections handled is indicated by the color blue and is shown behind the number of bypassed connections, which is indicated by the color green.

Figure 16-48 Video Connection Details Chart

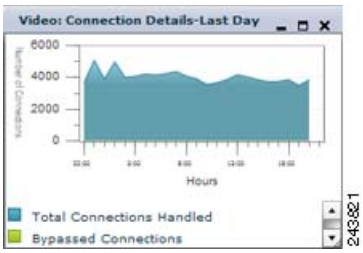

Video: Acceleration Bypass Reason

The Video Acceleration Bypass Reason pie chart (see Figure 16-49) displays the reasons that video traffic is not accelerated: Windows Media VOD, aggregate bitrate overload, other reasons, stream bitrate overload, session count overload, or unsupported transmission type (this means unsupported transport, unsupported player, or unsupported protocol).

Figure 16-49 Video Acceleration Bypass Reason Chart

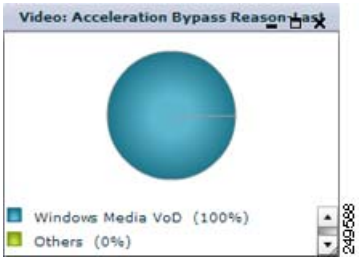

Video: Bandwidth Optimization

The Video Bandwidth Optimization chart (see Figure 16-50) displays the effective bandwidth capacity of the WAN link as a result of video acceleration, as a multiplier of its base capacity. The chart in the display is an area chart; data for all traffic is indicated by the color blue and is shown behind streaming video traffic, which is indicated by the color green.

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.

Figure 16-50 Video Bandwidth Optimization Chart

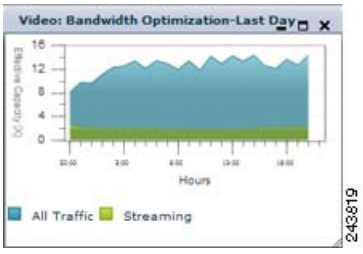

SSL

This section describes these charts:

- SSL: Acceleration Bypass Reason, page 16-33
- SSL: Connection Details, page 16-33
- SSL: Bandwidth Optimization, page 16-33

SSL: Acceleration Bypass Reason

The SSL Acceleration Bypass Reason pie chart (see Figure 16-51) displays the reasons that SSL traffic is not accelerated: version mismatch, unknown, non-matching domain, cipher mismatch, revocation failure, certificate verification failure, other failure, and non-SSL traffic.

Figure 16-51 SSL Acceleration Bypass Reason Chart

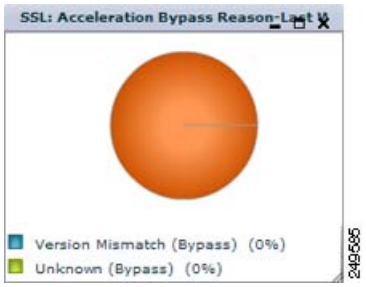

SSL: Connection Details

The SSL Connection Details chart (see Figure 16-52) displays the SSL session connection statistics, showing the total number of connections handled, the number of unaccelerated (bypassed) connections, and the number of dropped connections. The chart in the display is an area chart; the total number of connections handled is indicated by the color blue and is shown behind the number of bypassed connections in green, which is shown behind the number of dropped connections in orange.

Figure 16-52 SSL Connection Details Chart

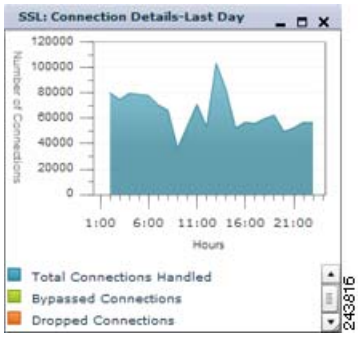

SSL: Bandwidth Optimization

The SSL Bandwidth Optimization chart (see Figure 16-53) displays the effective bandwidth capacity of the WAN link as a result of SSL acceleration, as a multiplier of its base capacity. The chart in the display is an area chart; data for all traffic is indicated by the color blue and is shown behind SSL traffic, which is indicated by the color green.

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.

Figure 16-53 SSL Bandwidth Optimization Chart

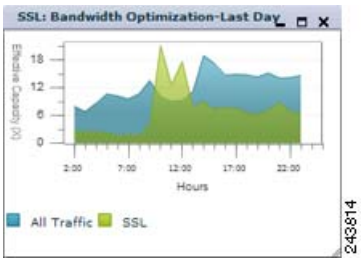

Platform Charts

This section describes these charts:

- Managed Devices Information, page 16-34
- CPU Utilization, page 16-34

Managed Devices Information

The Managed Devices Information chart (see Figure 16-54) displays information about the number of devices managed by the WAAS Central Manager, the device status, number of alarms, and software versions. You cannot customize the data shown in this chart and you cannot export the data, though you can print it. You can add this chart only to the Platform tab in the system dashboard.

Figure 16-54 Managed Devices Information Chart

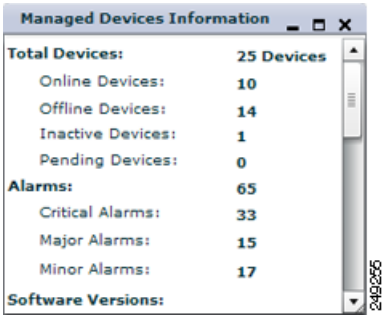

CPU Utilization

The CPU Utilization chart (see Figure 16-55) displays the percentage of CPU utilization for the device. This chart is available only when a specific WAAS device is selected. It is the only chart that can be added to the **Report > Manage Reports > CPU Usage** report page.

Figure 16-55 CPU Utilization Chart

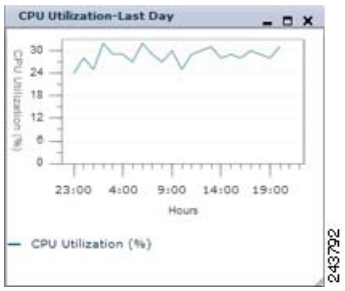

Using Predefined Reports to Monitor WAAS

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 drawer in the navigation pane. 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 Report drawer, as described in the "Viewing and Editing Reports" section on page 16-50.

The following predefined reports are available at the WAAS system level, the location level, and the WAE device level:

- Optimization
  - Traffic Summary Report, page 16-36
  - Optimization Summary Report, page 16-39
  - Optimization Details Report, page 16-39
- Acceleration
  - HTTP Acceleration Report, page 16-40
  - Video Acceleration Report, page 16-41
  - SSL Acceleration Report, page 16-42
  - MAPI Acceleration Report, page 16-42
  - NFS Acceleration Report, page 16-43

The following predefined reports are available only at the WAAS System level and the WAE device level:

- Topology Report, page 16-44

The following predefined reports are available only at the WAE device level:

- Optimization
  - Connections Statistics Report, page 16-45
- Acceleration
  - CIFS Acceleration Report, page 16-47
- Platform
  - CPU Statistics Report, page 16-47
  - Disks Report, page 16-48

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 PDF 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 3-14.

To view a location level report, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Locations**.

Step 2 Click the **Edit** icon next to the location for which you want to view reports.

Step 3 From the navigation pane, choose **Monitor** and choose the report from the Optimization or Acceleration menus.

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 16-51.

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 9-17.

Traffic Summary Report

The Traffic Summary report (see Figure 16-56) displays a summary of all traffic.

Figure 16-56 Traffic Summary Report

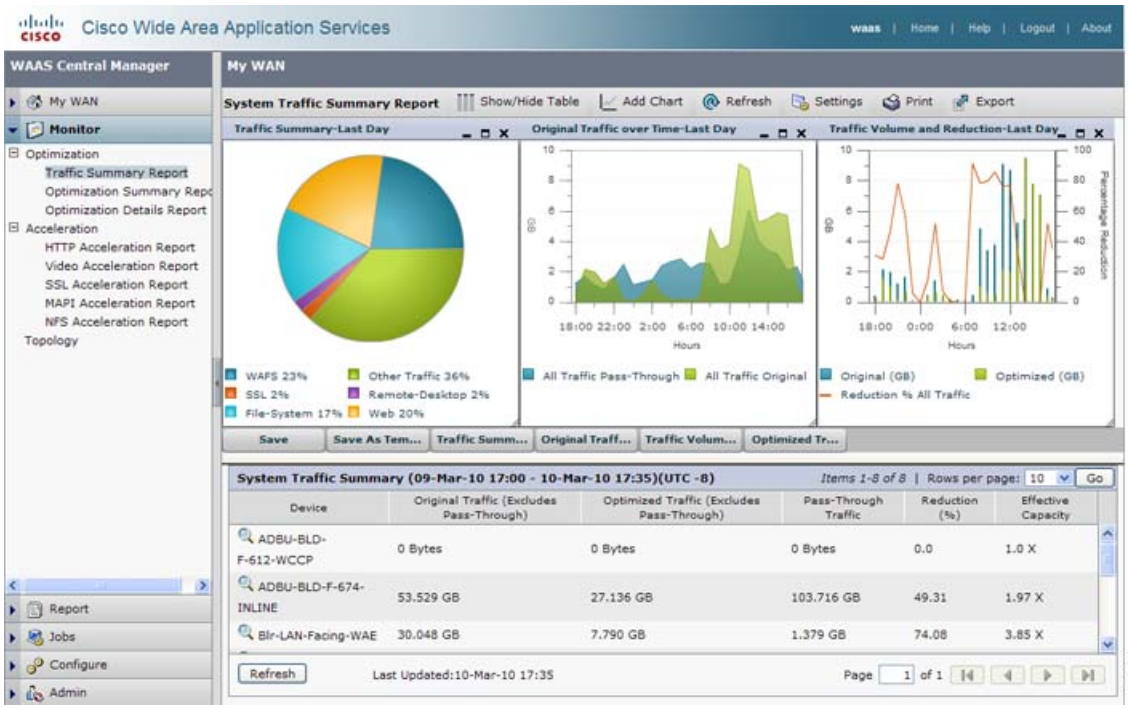

The following charts are included:

- Traffic Summary, page 16-14
- Original Traffic Over Time, page 16-15
- Traffic Volume and Reduction, page 16-17
- Optimized Traffic Over Time, page 16-17 (this chart is minimized)

The Traffic Summary table is displayed below the charts. 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 16-4.

You can sort the table by clicking on any column heading to sort on data in that column. A small triangle control appears below 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. The terms used in the table are defined as follows:

- Original Inbound—Traffic that is entering the WAE from the LAN (clients) and needs to be optimized before being sent out on the WAN to a peer WAE.
- Original Outbound—Traffic that is exiting the WAE to the LAN (clients) after being received on the WAN from a peer WAE.
- Optimized Inbound—Traffic that is entering the WAE 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 WAE to the WAN and a peer WAE after being optimized.
- Pass-Through—Traffic that is being passed through the WAE 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 values Reduction % (incl. pass-through), Reduction % (excl. pass-through), and Effective Capacity are calculated using these added values of all devices.

Table 16-4 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.)
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. 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 "Clustering Inline WAEs" section on page 4-50).[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; this 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.

   Say W1 and W2 are part of a serial cluster and W1 is towards 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 towards 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)

Optimization Summary Report

The Optimization Summary report (see Figure 16-57) displays the optimization summary.

Figure 16-57 Optimization Summary Report

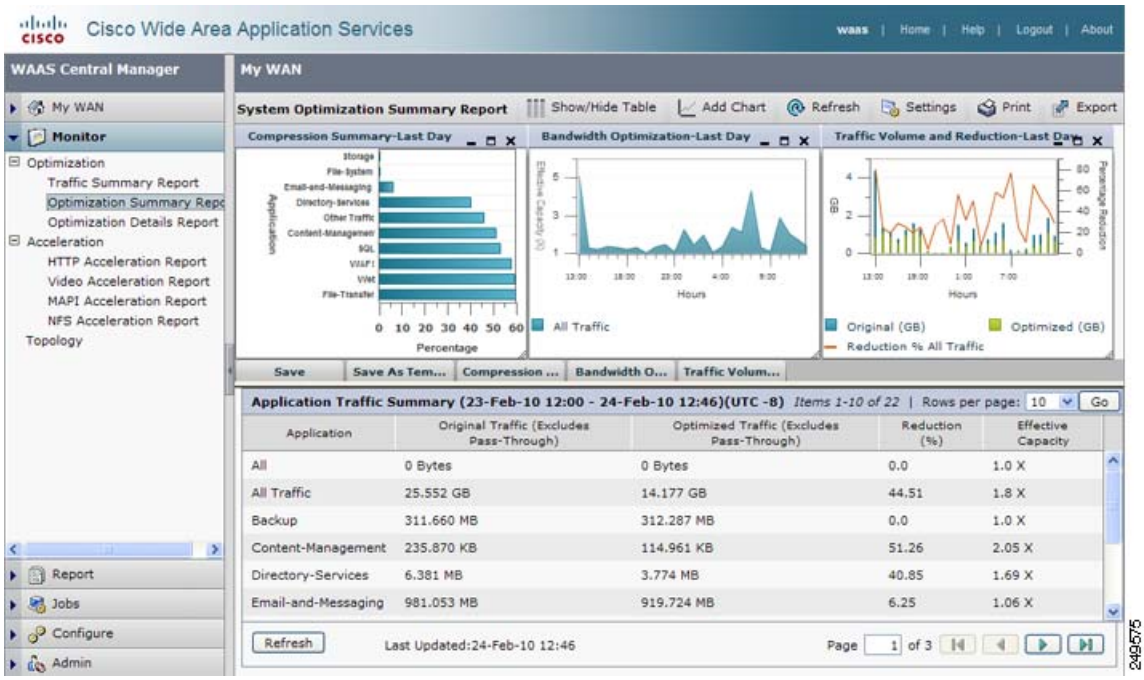

The following charts are included:

- Compression Summary, page 16-16
- Bandwidth Optimization, page 16-18
- Traffic Volume and Reduction, page 16-17

The Application Traffic Summary table is displayed below the charts. Each row in the table displays the total traffic information for each application. The column descriptions and formulas used to calculate the values are listed in Table 16-4. Only the following columns are included:

- Application
- Original Traffic (Excludes Pass-Through)
- Optimized Traffic (Excludes Pass-Through)
- Reduction (%) (this is excluding pass-through traffic)
- Effective Capacity (this is excluding pass-through traffic)

Optimization Details Report

The Optimization Details report (see Figure 16-58) displays the optimization details.

Figure 16-58 Optimization Details Report

The following charts are included:

- Compression Over Time, page 16-16
- Traffic Volume and Reduction, page 16-17
- Bandwidth Optimization, page 16-18

The Application Traffic Details table is displayed below the charts. Each row in the table displays the total traffic information for each application. The column descriptions and formulas used to calculate the values are listed in Table 16-4. Only the following columns are included:

- Application
- Original Traffic (Excludes Pass-Through)
- Optimized Traffic (Excludes Pass-Through)
- Pass-Through Traffic
- Reduction (%) (this is excluding pass-through traffic)
- Effective Capacity (this is excluding pass-through traffic)

HTTP Acceleration Report

The HTTP Acceleration report displays the HTTP acceleration statistics.

The following charts are included on two tabs:

- Summary tab:
  - HTTP: Estimated Time Savings, page 16-19
  - HTTP: Bandwidth Optimization, page 16-19

  - HTTP: Connection Details, page 16-19
- Details tab:
  - HTTP: Response Time Savings, page 16-20
  - HTTP: Optimization Count, page 16-20
  - HTTP: Optimization Techniques, page 16-21

The HTTP Acceleration Statistics table is displayed below the charts. Each row in the table displays statistical information for each device that is registered to this Central Manager. The data is described in Table 16-5.

Table 16-5 HTTP Acceleration Statistics Table

Table Column	Description and Formulas Used to Calculate Value
Device	The device name.
New Connections Handled	Reports the number of HTTP connections handled for the time period.
Active Connections	Reports the number of connections currently being handled by the HTTP accelerator.
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)

Video Acceleration Report

The Video Acceleration report displays the video acceleration statistics.

The following charts are included on two tabs:

- Summary tab:
  - Video: Stream Optimization, page 16-31
  - Video: Bandwidth Optimization, page 16-32
  - Video: Connection Details, page 16-31
- Details tab:
  - Video: Acceleration Bypass Reason, page 16-32

The Video Acceleration Statistics table is displayed below the charts. Each row in the table displays statistical information for each device that is registered to this Central Manager. The data is described in Table 16-6.

Table 16-6 Video Acceleration Statistics Table

Table Column	Description
Device	The device name.
New Connections Handled	Reports the number of video connections handled for the time period.
Active Connections	Reports the number of connections currently being handled by the video accelerator.
New Bypassed Connections	Reports the number of connections initially received by the video accelerator and then pushed down to the generic accelerator.

SSL Acceleration Report

The SSL Acceleration report displays the SSL acceleration statistics.

The following charts are included:

- SSL: Bandwidth Optimization, page 16-33
- SSL: Connection Details, page 16-33
- SSL: Acceleration Bypass Reason, page 16-33

The SSL Acceleration Statistics table is displayed below the charts. Each row in the table displays statistical information for each device that is registered to this Central Manager. The data is described in Table 16-7.

Table 16-7 SSL Acceleration Statistics Table

Table Column	Description
Device	The device name.
Handled Connections	Reports the number of SSL connections handled for the time period.
Active Connections	Reports the number of connections currently being 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.

MAPI Acceleration Report

The MAPI Acceleration report displays the MAPI acceleration statistics.

The following charts are included on two tabs:

- Summary tab:
  - MAPI: Estimated Time Savings, page 16-26
  - MAPI: Bandwidth Optimization, page 16-28
  - MAPI: Connection Details, page 16-27
- Details tab:
  - MAPI: Request Optimization, page 16-26

- MAPI: Response Time Optimization, page 16-26
- MAPI: Versions Detected, page 16-26
- MAPI: Acceleration Bypass Reason, page 16-27

The MAPI Acceleration Statistics table is displayed below the charts. Each row in the table displays statistical information for each device that is registered to this Central Manager. The data is described in Table 16-8.

Table 16-8 MAPI Acceleration Statistics Table

Table Column	Description and Formulas Used to Calculate Value
Device	The device name.
New Connections Handled	Reports the number of MAPI connections handled for the time period.
Active Connections	Reports the number of connections currently being handled by the MAPI accelerator.
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.
% 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))

NFS Acceleration Report

The NFS Acceleration report displays the NFS acceleration statistics.

The following charts are included on two tabs:

- Summary tab:
  - NFS: Estimated Time Savings, page 16-29
  - NFS: Bandwidth Optimization, page 16-30
  - NFS: Connection Details, page 16-30
- Details tab:
  - NFS: Request Optimization, page 16-28
  - NFS: Response Time Optimization, page 16-29
  - NFS: Versions Detected, page 16-29

– NFS: Acceleration Bypass Reason, page 16-30

The NFS Acceleration Statistics table is displayed below the charts. Each row in the table displays statistical information for each device that is registered to this Central Manager. The data is described in Table 16-9.

Table 16-9 NFS Acceleration Statistics Table

Table Column	Description and Formulas Used to Calculate Value
Device	The device name.
New Connections Handled	Reports the number of NFS connections handled for the time period.
Active Connections	Reports the number of connections currently being handled by the NFS accelerator.
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.
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))

Topology Report

The Topology report at the system level displays a topology map that shows a graphical representation of all the connections between the WAE 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 TFO 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 TFO 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.

Connections Statistics Report

The Connections Statistics report displays a connection summary table for the device and is available at the device level only. The table displays all of the TCP connections handled by the device and corresponds to the **show statistics connection** EXEC mode command. (See Figure 16-59.)

Figure 16-59 Device Connections Summary Table

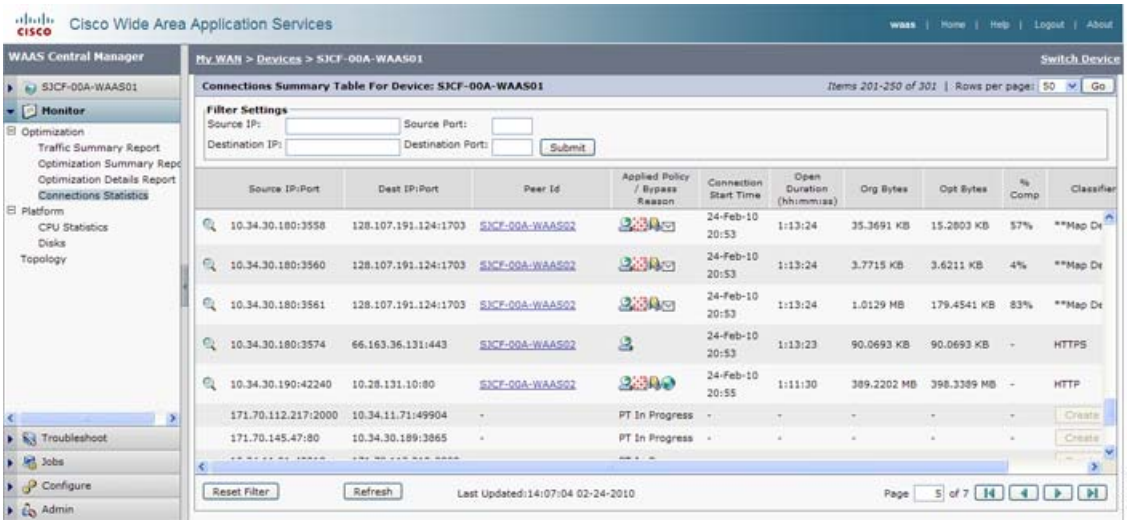

This window displays the following information about each connection:

- 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. (See Figure 16-60.)

Figure 16-60 Connection Details

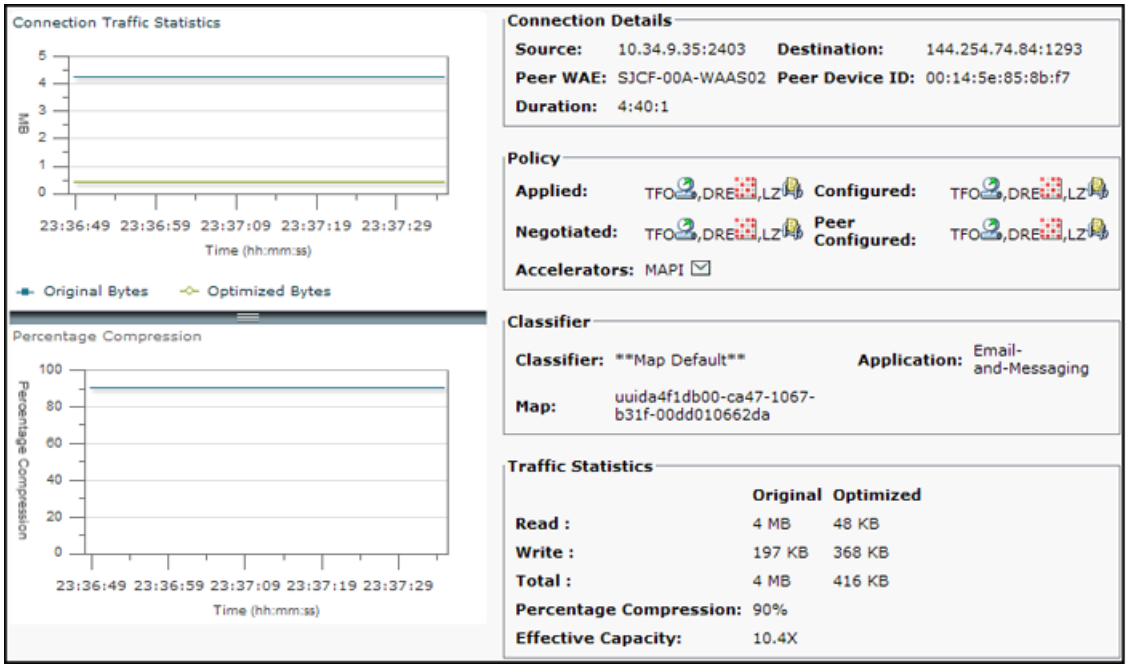

Note If the value for Percentage Compression is negative, the Percentage Compression and Effective Capacity values do not appear.

CIFS Acceleration Report

The CIFS Acceleration report displays the CIFS acceleration statistics and is available at the device level only. Different charts are available depending on the CIFS accelerator mode (transparent CIFS accelerator, legacy CIFS Edge, or legacy CIFS Core).

For a transparent CIFS accelerator device, the following charts are included on two tabs:

- Summary tab:
  - CIFS: Connection Statistics, page 16-23
  - CIFS: File Optimization, page 16-23
  - CIFS: Request Optimization, page 16-21
  - CIFS: Cache Utilization, page 16-22
- Details tab:
  - CIFS: Cached Objects, page 16-22
  - CIFS: Client Average Throughput, page 16-23

For a legacy CIFS Edge device, the following charts are included on two tabs:

- Summary tab:
  - CIFS: Connection Statistics, page 16-23
  - CIFS: File Optimization, page 16-23
  - CIFS: Request Optimization, page 16-21
  - CIFS: Cache Utilization, page 16-22
- Details tab:
  - CIFS: Cached Objects, page 16-22
  - CIFS: Connected CIFS Core Count, page 16-24
  - CIFS: CIFS Edge-CIFS Core Traffic, page 16-24
  - CIFS: Client Average Throughput, page 16-23

For a legacy CIFS Core device, the following charts are included:

- Summary tab:
  - CIFS: Connected CIFS Edge Count, page 16-25
  - CIFS: CIFS Core Traffic, page 16-25

Note When you use the **Print** icon in the taskbar to print the CIFS Acceleration report to a PDF 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.

CPU Statistics Report

The CPU Statistics report displays the CPU Utilization chart.

To change the report time frame, click **Settings**.

Disks Report

The Disks Report displays physical and logical disk information. (See Figure 16-61.)

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.

Figure 16-61 Disk Information for Device Window

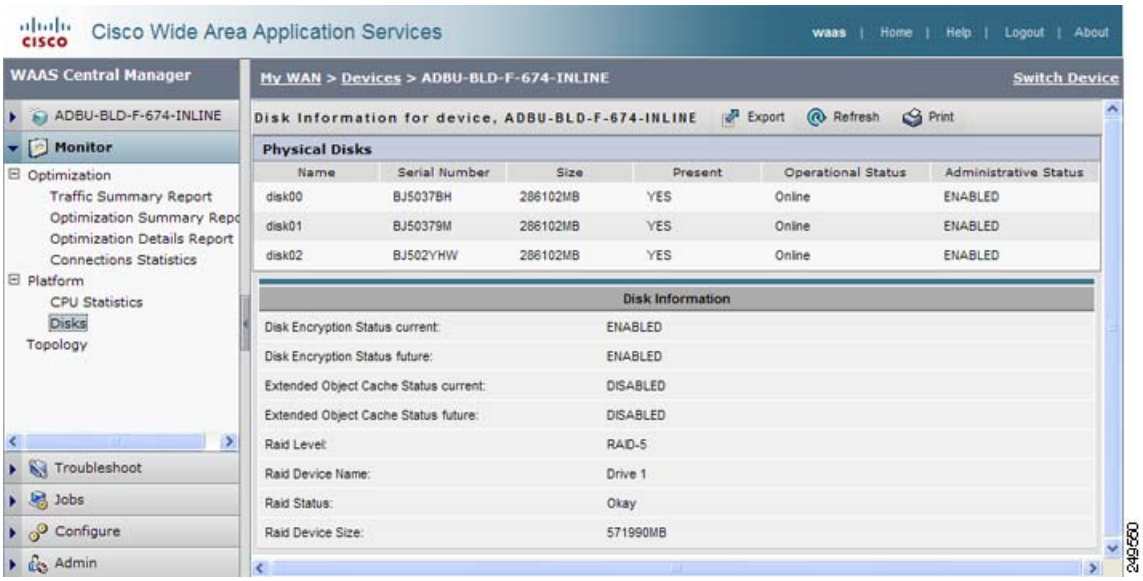

Managing Reports

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 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 16-49
- Viewing and Editing Reports, page 16-50
- Scheduling Reports, page 16-51
- Managing Scheduled Reports, page 16-52

Creating Custom Reports

A report consists of one or more charts in the chart panel (upper portion of the window) and a table in the table panel (lower portion of the window). The system and device dashboard displays are examples of predefined reports, along with the other reports available in the Monitor drawer.

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 GUI navigation pane, choose **Report > Manage Reports**.

Step 2 Click the **Create New Report** icon in the taskbar. The Custom Report window appears, as shown in Figure 16-62.

Figure 16-62 Creating a Report

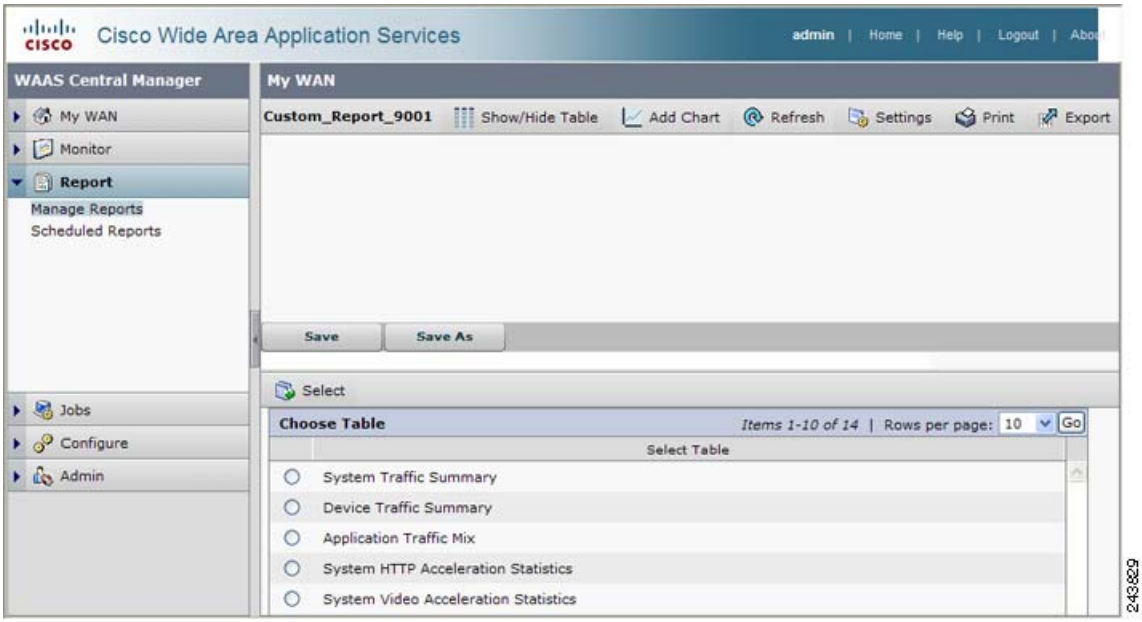

Step 3 Choose a table to display in the lower table panel of the report by clicking the radio button next to the table in the Choose Table area and then clicking the **Select** button above the Choose Table area.

Step 4 Add one or more charts to be displayed in the upper chart panel of the report by clicking the **Add Chart** icon in the taskbar. The Add Chart window is displayed, as shown in Figure 16-7.

Step 5 Expand any of the chart categories by clicking on the plus sign next to the category.

Step 6 Check the box next to each chart that you want to display. See the "Chart Descriptions" section on page 16-14 for a description of the charts.

A report can contain a maximum of six charts.

Step 7 Click **Add**.

Step 8 Customize any of the chart settings by clicking the **Settings** icon in the taskbar. For more information, see the "Configuring Chart Settings" section on page 16-12.

Step 9 Click the **Save As** button below the chart panel to save the report with a new name. The Save As popup window appears.

(If you use the Save button, the report is saved with a default name such as Custom_Report_9001.)

Step 10 Enter the report name and any notes about the report.

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

Step 11 Click **Submit**.

If you want to delete a chart from a dashboard or report, click the **Close** button on the chart and save the report.

All of the taskbar icons above the charts panel and the buttons below the charts panel work as described in the "Customizing a Dashboard or Report" section on page 16-10.

Viewing and Editing Reports

To view or edit a report, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Report > Manage Reports**.

Step 2 Check the box next to 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 navigate to another page of the Reports table.

Step 3 To delete a report, click the **Delete** icon in the taskbar.

Step 4 To view or edit a report, click the **Edit** icon in the taskbar. The report is displayed.

Alternatively, click the magnifying glass icon next to a report as a shortcut to viewing it.

Step 5 If you want to change any of the charts in the report, use the standard chart editing methods as described in the "Creating Custom Reports" section on page 16-49.

Step 6 Click **Save** or **Save As** to save the report.

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

When editing a report, you cannot change the table that is displayed in the table panel. If you want to change the table, you must create a new report.

When viewing a report, click the **Refresh** button below the table panel to refresh the data in the table.

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 PDF copy of the report can be e-mailed.

To schedule a report, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Report > Manage Reports**.

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 navigate to another page of the Reports table.

Step 3 Click the **Schedule** icon in the taskbar. The scheduling window appears, as shown in Figure 16-63.

Figure 16-63 Scheduling a Report

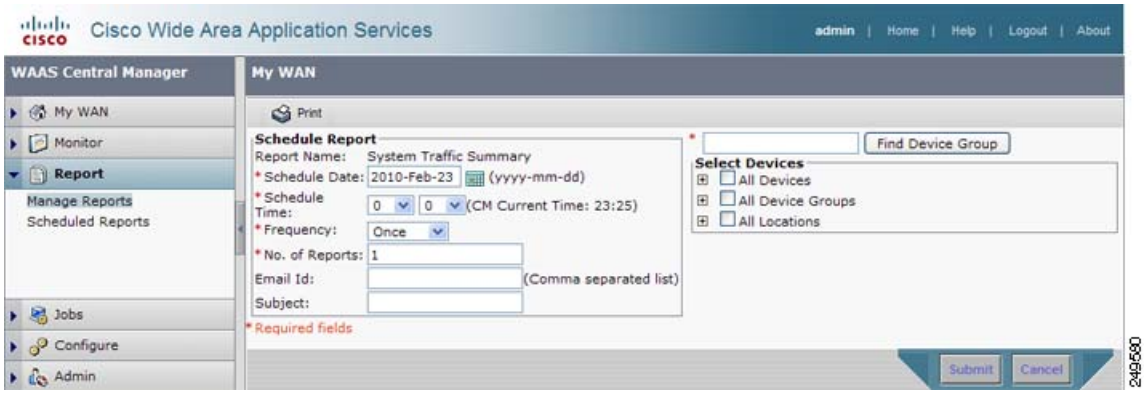

Step 4 In the Schedule Date field, enter the schedule date in the format YYYY-MM-DD, or click the calendar icon to display a calendar popup window from which to choose the date.

Step 5 In the Schedule Time field, choose the hours and minutes from the drop-down lists. The time represents the local time at the WAAS Central Manager.

Step 6 In the Frequency drop-down list, choose Once, Daily, Weekly, or Monthly for the report frequency.

Step 7 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 8 In the Email Id field, enter the e-mail addresses of the report recipients, separated by commas.

Step 9 In the Subject field, enter the subject of the e-mail message.

Step 10 In the Select Devices area, select the devices that are to be included in the statistics for the report. Place a check in the box next to each device, location, or device group that you want to include. Individual devices are listed only in the All Devices list.

To find (highlight) a device group in a long list, enter the device group name in the field above the list and click **Find Device Group**. The search is case sensitive.

Step 11 Click **Submit**.

Step 12 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 9-24.

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 GUI navigation pane, choose **Report > Scheduled Reports**. The Scheduled Reports window lists the scheduled reports.

Step 2 (Optional) If you want to view a report, click the plus sign on the row for the schedule. The row expands to show the report instances for the schedule. Each instance shows the report name, devices selected, device groups selected, locations selected, scheduled time, completed time, frequency, status, number of pending reports remaining (for reoccurring reports), and the user that scheduled the report (this last column is visible only to admin users).

Check the box next to the report instance that you want to view and click the **View Completed Report** icon on the taskbar; or, click the Completed link in the status column. You can view only completed reports. If a report instance has a status of Not Started or In Progress, you must wait for it to complete before you can view the report.

Note For each completed instance of a scheduled report, the Frequency column will show Once and the Completed Time will show the date and time the report was generated. These instances are the viewable instances of the report—one for each instance. There is also one instance shown with a Frequency of Daily, Weekly, or Monthly (depending on the schedule period), and no Completed Time. That instance is not viewable and represents the scheduled report object.

Step 3 (Optional) If you want to delete a report, check the box next to one or more report instances that you want to delete and click the **Delete Selected Reports** icon on the taskbar. You must select report instances, not the schedule name above a group of instances. If all report instances are deleted, the schedule is also deleted.

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 9-17.

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 will see the changes, and other users (including admin users) will continue to see the report with default settings.

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
- Example Using NetQoS for Flow Monitoring

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 16-55.)

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 16-54.) 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 **My WAN > Manage Device Groups > 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 In the navigation pane, 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 16-10 describes the four different alarms that may be raised when errors occur with flow monitoring.

Table 16-10 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 dependant 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, 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 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group for which you want to enable system logging.

Step 3 From the navigation pane, choose **Configure > Monitoring > Log Settings > System Log**. The System Log Settings window appears. (See Figure 16-64.)

Figure 16-64 System Log Settings Window

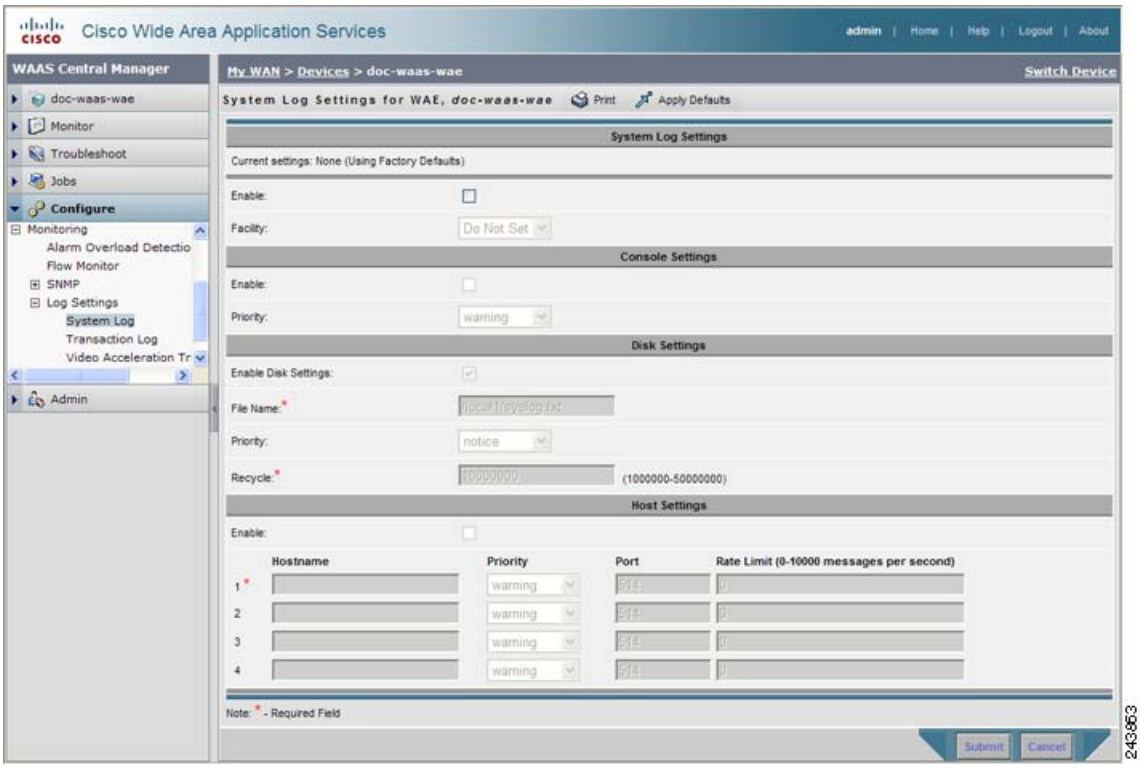

Step 4 Under the System Log Settings section, check the **Enable** check box to enable system logging. By default, this option is disabled.

Step 5 From the Facility drop-down list, choose the appropriate facility.

Step 6 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 16-11 on page 16-57 for a list of priority levels.)

Step 7 Enable syslog files to be sent to 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 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 16-11 on page 16-57 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 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 8 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 16-58."

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 16-11 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 9 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, page 16-57
- Multiple Hosts for System Logging, page 16-58

Priority Levels

Table 16-11 lists the different priority levels of detail to send to the recipient of the syslog messages for a corresponding event.

Table 16-11 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.
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, page 16-58
- Transaction Logs, page 16-60

Enabling Transaction Logging

To enable transaction logging for TFO flows and video streams, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the device or device group for which you want to enable system logging. The Device Dashboard window or the Modifying Device Group window appears.

Step 3 From the navigation pane, 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 16-65.) The Video Transaction Log Settings window looks the same, but does not include the General Settings area at the top.

Figure 16-65 Transaction Log Settings Window

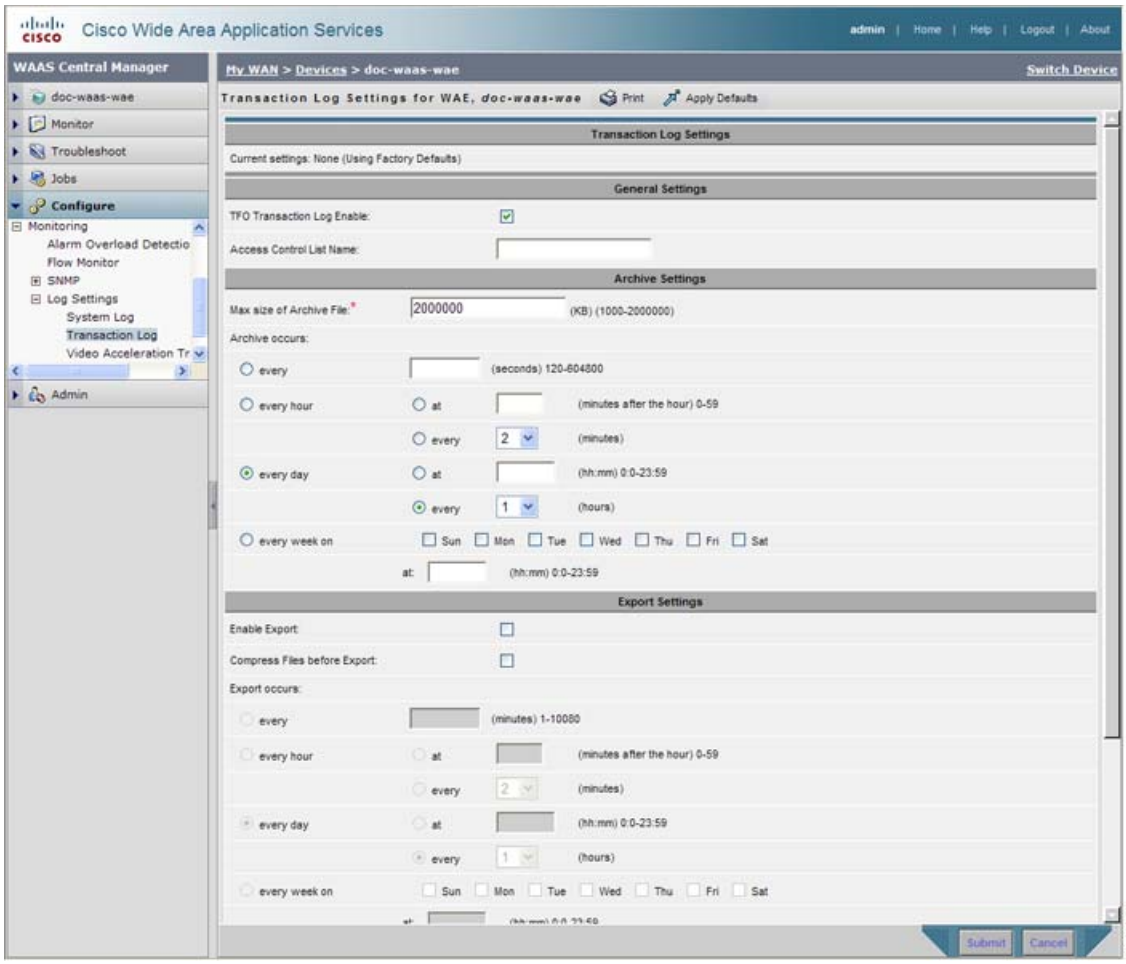

Step 4 Under the General Settings heading, check the **TFO Transaction Log Enable** check box to enable transaction logging. This check box does not appear for video transaction logging.

The fields on the window become active.

Step 5 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 6 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 7 Configure the fields in the Export Settings section to export the transaction log file to an FTP server.

Table 16-12 describes the fields in the Export Settings section.

Table 16-12 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. • 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, place a check in the SFTP check box.

Step 8 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 GUI navigation pane, choose **Admin > Logs > System Messages**. The System Message Log window appears. (See Figure 16-66.)

Figure 16-66 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 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 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 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 GUI navigation pane, choose **Admin > Logs > Audit Trail Logs**.

The Audit Log window appears. (See Figure 16-67.) 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 16-67 Audit Log Window

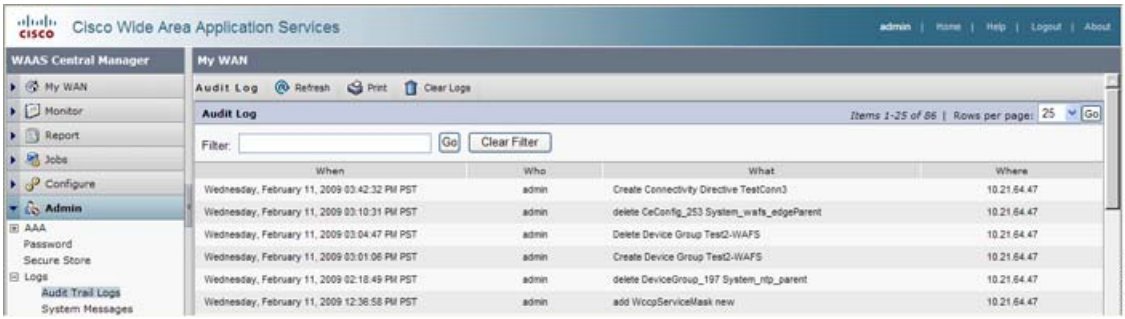

Step 2 Choose a number from the Rows drop-down list to determine the number of rows that you want to display.

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 16-61.

To view the logged information for a WAAS device, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**. The Devices window appears.

Step 2 Click the **Edit** icon next to the device for which you want to view the system message log details. The Device Dashboard window appears.

Step 3 In the navigation pane, choose **Admin > Logs**. The System Message Log for Device window appears.

Step 4 Choose the type of messages to be displayed from the System Message Log drop-down list.

You can view the following types of messages in the system log:

- All (default)
- CLI
- Critical
- Database

Step 5 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 6 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.

Enabling the Kernel Debugger

The WAAS Central Manager GUI allows you to enable or disable access to the kernel debugger (kdb). Once enabled, kernel debugger is automatically activated when kernel problems occur.

To enable the kernel debugger, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the WAAS device (or device group) that you want to debug.

Step 3 In the navigation Pane, choose **Troubleshoot > Kernel Debugger**. The Kernel Debugger window appears.

Step 4 Check the **Enable** check box to enable the kernel debugger, and click **Submit**. By default, this option is disabled.

Troubleshooting Using Diagnostic Tests

WAAS includes various troubleshooting tools as described in the following sections:

- Troubleshooting Using the GUI, page 16-64
- Troubleshooting Using the CLI, page 16-65

Troubleshooting 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**).

Step 2 Click the **Edit** icon next to the name of the device (or device group) for which you want to perform diagnostic tests.

Step 3 In the navigation pane, choose **Troubleshoot > Diagnostics Tests**. The Diagnostic Tool window appears.

Step 4 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 device status, presence of coredump files or alarms of major or critical severity.
- Basic Configuration—Check device basic network configuration.
- Basic Connectivity—Check device connectivity to configured external devices (DNS, authentication, NTP servers, etc.).
- Physical Interface—Check configuration and operation of device physical interfaces.
- Configuration Security—Check running configuration for potentially malicious (XSS) entries.
- Traffic Optimization—Check TFO configuration and operation.
- WCCP configuration and operation—Check configuration and operation of WCCP traffic interception.
- Inline configuration and operation—Check configuration and operation of inline group interfaces.
- WAFS configuration and operation—Check configuration and operation of WAFS services.

Step 5 Click **Run**.

Step 6 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.

Troubleshooting 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, then 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, then the output can be filtered based on 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 GUI navigation pane, choose **My WAN > Manage Devices**.

Step 2 Click the **Edit** icon next to the device for which you want to issue a **show** or **clear** command.

Step 3 From the navigation pane, choose **Troubleshoot > CLI Commands** and then click either **Show Commands** or **Clear Commands**.

Step 4 From the drop-down list, choose a **show** or **clear** command.

Step 5 Enter arguments for the command, if any.

Step 6 Click **Submit** to display the command output.

A window appears, displaying the command output for that device.

You can also use the **show** EXEC commands from the CLI. For more information, see the *Cisco Wide Area Application Services Command Reference*.

CHAPTER 17

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, page 17-1
- Checklist for Configuring SNMP, page 17-8
- Preparing for SNMP Monitoring, page 17-9
- Enabling SNMP Traps, page 17-10
- Defining SNMP Traps, page 17-12
- Specifying the SNMP Host, page 17-14
- Specifying the SNMP Community String, page 17-15
- Creating SNMP Views, page 17-16
- Creating an SNMP Group, page 17-17
- Creating an SNMP User, page 17-18
- Configuring SNMP Asset Tag Settings, page 17-20
- Configuring SNMP Contact Settings, page 17-20
- Configuring SNMP Trap Source Settings, page 17-21

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, page 17-2
- Supported SNMP Versions, page 17-3
- SNMP Security Models and Security Levels, page 17-3
- Supported MIBs, page 17-4
- Downloading MIB Files, page 17-8
- Enabling the SNMP Agent on a WAAS Device, page 17-8

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 17-10.

3. After locating the specified information in the MIB, the agent uses SNMP to send the information to the SNMP management station.

Figure 17-1 illustrates these SNMP operations for an individual WAAS device.

Figure 17-1 SNMP Components in a WAAS Network

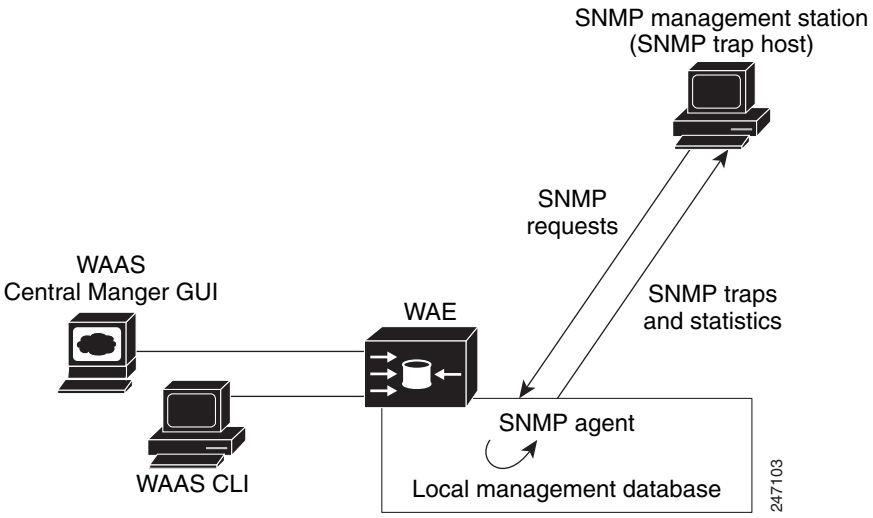

Supported SNMP Versions

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 17-1 describes the combinations of security models and security levels.

Table 17-1 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:

- ACTONA-ACTASTOR-MIB
- CISCO-CDP-MIB
- CISCO-CONFIG-MAN-MIB
- CISCO-CONTENT-ENGINE-MIB
- CISCO-ENTITY-ASSET-MIB
- CISCO-SMI

- ENTITY-MIB
- EVENT-MIB
- HOST-RESOURCES-MIB
- MIB-II
- SNMP-COMMUNITY-MIB
- SNMP-FRAMEWORK-MIB
- SNMP-NOTIFICATION-MIB
- SNMP-TARGET-MIB
- SNMP-USM-MIB
- SNMPv2-MIB
- SNMP-VACM-MIB

ACTONA-ACTASTOR-MIB

This MIB provides statistics for the CIFS transparent accelerator and statistics and log traps for the legacy mode WAFS component in WAAS.

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.

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

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:

- ifInUcastPkts
- ifInUnknownProtos
- ifOutUcastPkts
- 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 17-18.

Checklist for Configuring SNMP

Table 17-2 describes the process for enabling SNMP monitoring on a WAAS device or device group.

Table 17-2 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 17-9.
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 17-10. To define additional traps, see the "Defining SNMP Traps" section on page 17-12.

Table 17-2 Checklist for Configuring SNMP (continued)

Task	Additional Information and Instructions
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 17-14.
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 17-15.
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 17-16.
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 17-17.
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 17-18.
8. Configure SNMP contact settings.	For more information, see the "Configuring SNMP Contact Settings" section on page 17-20.

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**). The Devices or Device Groups window appears depending on your selection.

Step 2 Click the **Edit** icon next to the device or device group that you want to configure SNMP traps for. The Device Dashboard window appears.

Step 3 In the navigation pane, choose **Configure > Monitoring > SNMP > General Settings**. The SNMP General Settings window appears. (See Figure 17-2.) Table 17-3 describes the fields in this window.

Figure 17-2 SNMP General Settings Window

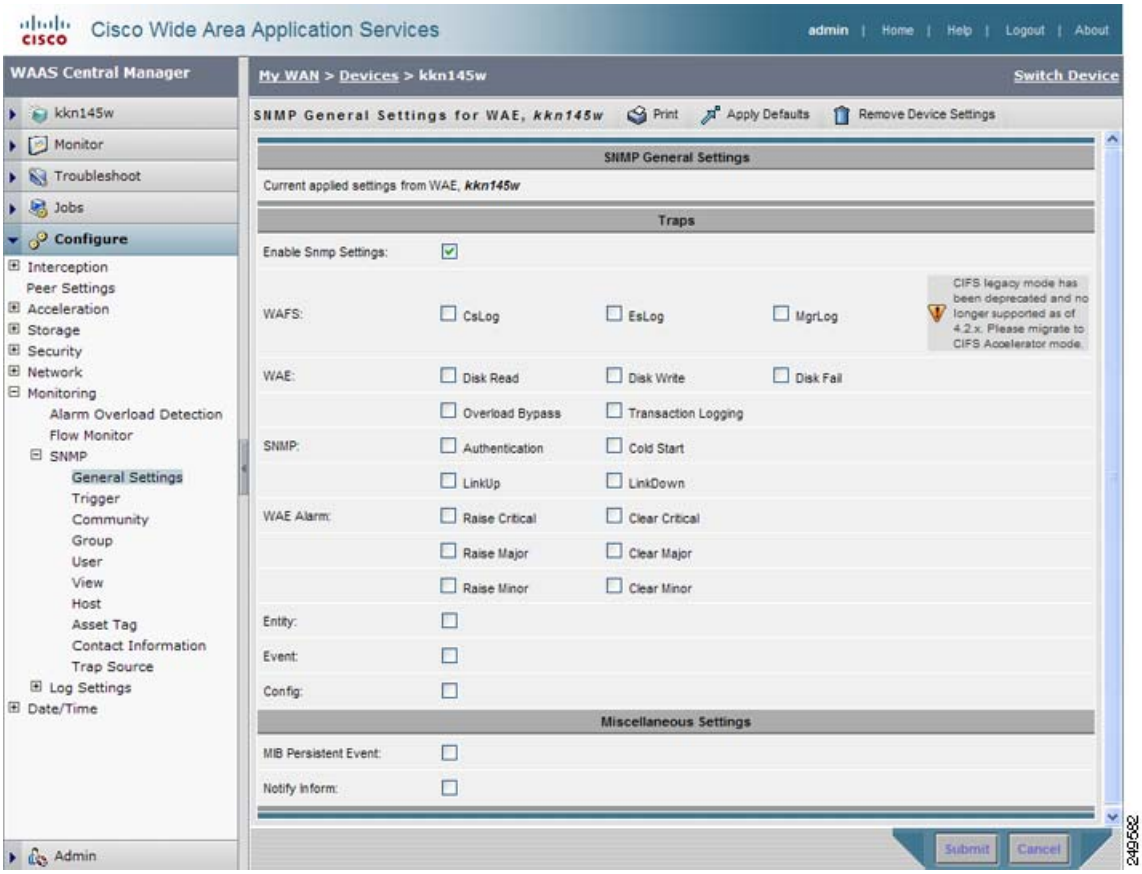

Table 17-3 SNMP General Settings

GUI Parameter	Function
Traps	
Enable Snmp Settings	Enables SNMP traps.

Table 17-3 SNMP General Settings (continued)

GUI Parameter	Function
WAFS	Enables SNMP WAFS traps: • CsLog—Enables WAFS legacy mode Core Server error log traps. • EsLog— Enables WAFS legacy mode Edge Server error log traps. • MgrLog—Enables WAAS Central Manager error log 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.
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 4 Check the appropriate check boxes to enable SNMP traps.

Step 5 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 define additional SNMP traps for other MIB objects of interest to your particular configuration, see the "Defining SNMP Traps" section on page 17-12.

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**). The Devices or Device Groups window appears depending on your selection.

Step 2 Click the **Edit** icon next to the device or device group that you want to define SNMP traps for. The Device Dashboard window appears.

Step 3 In the navigation pane, 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 17-4.

Step 4 In the taskbar, click the **Create New SNMP Trigger List Entry** icon. The Creating New SNMP Trigger window appears. Table 17-4 describes the fields in this window.

Table 17-4 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 17-4 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 5 In the appropriate fields, enter the MIB name, frequency, test, sample type, threshold value, and comments.

Step 6 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.

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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**). The Devices or Device Groups window appears.

Step 2 Click the **Edit** icon next to the device or device group for which you want to define an SNMP host. The Device Dashboard window or the Modifying Device Groups window appears.

Step 3 In the navigation pane, choose **Configure > Monitoring > SNMP > Host**. The SNMP Hosts window appears.

Step 4 In the taskbar, click the **Create New SNMP Host** icon. The Creating New SNMP Host window appears. Table 17-5 describes the fields in this window.

Table 17-5 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.

Table 17-5 SNMP Host Settings (continued)

GUI Parameter	Function
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 5 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 6 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**). The Devices or Device Groups window appears.

Step 2 Click the **Edit** icon next to the device or device group for which you want to configure an SNMP community setting.

Step 3 In the navigation pane, choose **Configure > Monitoring > SNMP > Community**. The SNMP Community Strings window appears.

Step 4 In the taskbar, click the **Create New SNMP Community String** icon. The Creating New SNMP Community String window appears. Table 17-6 describes the fields in this window.

Table 17-6 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 5 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 6 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**). The Devices or Device Groups window appears.

Step 2 Click the **Edit** icon next to the device or device group for which you want to create an SNMPv2 view.

Step 3 In the navigation pane, choose **Configure > Monitoring > SNMP > View**. The SNMP Views window appears.

Step 4 In the taskbar, click the **Create New View** icon. The Creating New SNMP View window appears. Table 17-7 describes the fields in this window.

Table 17-7 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 5 In the appropriate fields, enter the view name, the family name, and the view type.

Step 6 Click **Submit**.

Step 7 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**). The Devices or Device Groups window appears.

Step 2 Click the **Edit** icon next to the device or device group for which you want to create an SNMP group. The Device Dashboard or the Modifying Device Group window appears.

Step 3 In the navigation pane, choose **Configure > Monitoring > SNMP > Group**. The SNMP Group Strings for WAE window appears.

Step 4 In the taskbar, click the **Create New SNMP Group String** icon. The Creating New SNMP Group String for WAE window appears. Table 17-8 describes the fields in this window.

Table 17-8 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 17-16.
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 17-16.
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 17-16.

Step 5 In the appropriate fields, enter the SNMP group configuration name, the security model, and the names of the read, write, and notify views.

Step 6 Click **Submit**.

Step 7 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**). The Devices or Device Groups window appears.

Step 2 Click the **Edit** icon next to the device or device group for which you want to create an SNMP user.

Step 3 In the navigation pane, choose **Configure > Monitoring > SNMP > User**. A list of SNMP users for the device or device group appears.

Step 4 In the taskbar, click the **Create New SNMP User** icon. The Creating New SNMP User window appears. Table 17-9 describes the fields in this window.

Table 17-9 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. 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.
Confirmation Password	Private password for confirmation. The reentered password must be the same as the one entered in the previous field.

Step 5 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 6 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**). The Devices or Device Groups window appears.

Step 2 Click the **Edit** icon next to the device or device group for which you want to define an SNMP asset tag. The Device Dashboard or the Modifying Device Groups window appears.

Step 3 In the navigation pane, choose **Configure > Monitoring > SNMP > Asset Tag**. The SNMP Asset Tag Settings window appears.

Step 4 In the Asset Tag Name field, enter a name for the asset tag.

Step 5 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 GUI navigation pane, choose **My WAN > Manage Devices** (or **Manage Device Groups**). The Devices or Device Groups window appears.

Step 2 Click the **Edit** icon next to the device or device group for which you want to configure an SNMP contact. The Device Dashboard or the Modifying Device Groups window appears.

Step 3 In the navigation pane, choose **Configure > Monitoring > SNMP > Contact Information**. The SNMP Contact Settings window appears.

Step 4 Enter a contact name and location in the provided fields.

Step 5 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 GUI navigation pane, choose **My WAN > Manage Devices**. The Devices window appears. (This setting is not supported from device groups.)

Step 2 Click the **Edit** icon next to the device or device group for which you want to configure an SNMP trap source. The Device Dashboard or the Modifying Device Groups window appears.

Step 3 In the navigation pane, choose **Configure > Monitoring > SNMP > Trap Source**. The SNMP Trap Source Settings window appears.

Step 4 From the Trap Source drop-down list, choose the interface to be used as the trap source. From the available gigabit Ethernet, standby, and port channel interfaces, only those with 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 5 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 Application Policies

The WAAS software includes over 150 predefined application policies 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 classifiers that WAAS will either optimize or pass through based on the policies that are provided with the system.

Before you create an application policy, we recommend that you review the predefined policies and modify them as appropriate. Often, you can more easily modify an existing policy than create a new one.

When reviewing Table A-1, note the following information:

- The subheadings represent the application names, and the associated classifiers are listed under these subheadings. For example, Authentication is a type of application and Kerberos is a classifier for that application.
- Applications with the word (*monitored*) next to them are monitored by the WAAS Central Manager, which can display statistics for only 20 applications 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 on a WAAS device. For more information, see the *Cisco Wide Area Application Services Command Reference*.
  - Modify the application settings so the WAAS Central Manager GUI displays statistics for the desired application. For more information, see Chapter 12, "Configuring Application Acceleration."

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.
- RE (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. RE operates on significantly larger streams and maintains a much larger compression history than LZ compression.
- LZ (compression)—Another compression technology that operates on smaller data streams and keeps limited compression history compared to RE.
- Application accelerator—A collection of individual application accelerators for the following traffic types: CIFS, EPM, HTTP, MAPI, NFS, SSL, and streaming video.

Table A-1 Predefined Traffic Policies

Classifier	WAAS Action	Destination Ports
Authentication		
Kerberos	Passthrough	88, 2053, 754, 888, 543, 464, 544, 749
SASL	Passthrough	3659
TACACS	Passthrough	49
Backup *(monitored)*		
Amanda	TFO	10080
BackupExpress	TFO	6123
CommVault	TFO	8400–8403
Connected-DataProtector	TFO	16384
IBM-TSM	LZ+TFO+DRE	1500-1502
Legato-NetWorker	TFO	7937, 7938, 7939
Legato-RepliStor	TFO	7144, 7145
Veritas-BackupExec	TFO	6101, 6102, 6106, 3527, 1125
Veritas-NetBackup	TFO	13720, 13721, 13782, 13785
CAD		
PDMWorks	LZ+TFO+DRE	30000, 40000
Call-Management		
Cisco-CallManager	Passthrough	2443, 2748
SIP-secure	Passthrough	5061
VoIP-Control	Passthrough	1300, 2428, 2000–2002, 1718–1720, 5060, 11000–11999
Conferencing		
CU-SeeMe	Passthrough	7640, 7642, 7648, 7649
ezMeeting	Passthrough	10101–10103, 26260–26261
Intel-Proshare	Passthrough	5713–5717
MS-NetMeeting	Passthrough	522, 1503, 1731
VocalTec	Passthrough	1490, 6670, 25793, 22555
Console		
SSL-Shell	Passthrough	614
Telnet	Passthrough	23, 107, 513
Telnets	Passthrough	992
Unix-Remote-Execution	Passthrough	514, 512
Content-Management *(monitored)*		
Documentum	LZ+TFO+DRE	1489
Filenet	LZ+TFO+DRE	32768–32774
ProjectWise-FileTransfer	LZ+TFO+DRE	5800

Table A-1 Predefined Traffic Policies (continued)

Classifier	WAAS Action	Destination Ports
Directory-Services *(monitored)*		
LDAP	LZ+TFO+DRE	389, 8404
LDAP-Global-Catalog	LZ+TFO+DRE	3268
LDAP-Global-Catalog-Secure	Passthrough	3269
LDAP-secure	Passthrough	636
Email-and-Messaging *(monitored)*		
HP-OpenMail	LZ+TFO+DRE	5755, 5757, 5766, 5767, 5768, 5729
Internet-Mail	LZ+TFO+DRE	25, 110, 143, 220
Internet-Mail-secure	TFO	995, 993, 465
Lotus-Notes	LZ+TFO+DRE	1352
MAPI[1]	LZ+TFO+DRE+ MAPI accelerator	UUID:a4f1db00-ca47-1067-b31f-00dd0106 62da
MDaemon	LZ+TFO+DRE	3000, 3001
NNTP	LZ+TFO+DRE	119
NNTP-secure	TFO	563
Novell-Groupwise	LZ+TFO+DRE	1677, 1099, 9850, 7205, 3800, 7100, 7180, 7101, 7181, 2800
PCMail-Server	LZ+TFO+DRE	158
QMTP	LZ+TFO+DRE	209
X400	LZ+TFO+DRE	102
Enterprise-Applications *(monitored)*		
SAP	LZ+TFO+DRE	3200-3219, 3221-3224, 3226-3267, 3270-3282, 3284-3305, 3307-3388, 3390-3399, 3600-3659, 3662-3699
Siebel	LZ+TFO+DRE	8448, 2320, 2321
File-System *(monitored)*		
AFS	LZ+TFO+DRE	7000–7009
Apple-AFP	LZ+TFO+DRE	548
NFS	LZ+TFO+DRE+ NFS accelerator	2049
Novell-NetWare	LZ+TFO+DRE	524
Sun-RPC	Passthrough	111
File-Transfer *(monitored)*		
BFTP	LZ+TFO+DRE	152
FTP-Control[2]	Passthrough	21
FTP-Data[2]	LZ+TFO+DRE	src20
FTPS[2]	TFO	990

Table A-1 Predefined Traffic Policies (continued)

Classifier	WAAS Action	Destination Ports
FTP-Control[2]	Passthrough	src989
Simple-FTP	LZ+TFO+DRE	115
TFTP	LZ+TFO+DRE	69
TFTPS	TFO	3713
Instant Messaging		
AOL	Passthrough	5190–5193
Apple-iChat	Passthrough	5297, 5298
IRC	Passthrough	531, 6660–6669
Jabber	Passthrough	5222, 5269
Lotus-Sametime-Connect	Passthrough	1533
MS-Chat	Passthrough	6665, 6667
MSN-Messenger	Passthrough	1863, 6891–6900
Yahoo-Messenger	Passthrough	5000, 5001, 5050, 5100
Name Services		
DNS	Passthrough	53
iSNS	Passthrough	3205
Service-Location	Passthrough	427
WINS	Passthrough	42, 137, 1512
Other *(monitored)*		
Basic-TCP-services	Passthrough	1–19
BGP	LZ+TFO+DRE	179
MS-EndPointMapper	EPM accelerator	135
MS-Message-Queuing	LZ+TFO+DRE	1801, 2101, 2103, 2105
NTP	Passthrough	123
Other-Secure	Passthrough	261, 448, 684, 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, 5989, 5990, 7674, 9802, 12109
SOAP	LZ+TFO+DRE	7627
Symantec-AntiVirus	LZ+TFO+DRE	2847, 2848, 2967, 2968, 38037, 38292
P2P *(monitored)*		
BitTorrent	Passthrough	6881–6889, 6969
eDonkey	Passthrough	4661, 4662
Gnutella	Passthrough	6346–6349, 6355, 5634
Grouper	Passthrough	8038

Table A-1 Predefined Traffic Policies (continued)

Classifier	WAAS Action	Destination Ports
HotLine	Passthrough	5500–5503
Kazaa	Passthrough	1214
Laplink-ShareDirect	Passthrough	2705
Napster	Passthrough	8875, 7777, 6700, 6666, 6677, 6688
Qnext	Passthrough	44, 5555
SoulSeek	Passthrough	2234, 5534
WASTE	Passthrough	1337
WinMX	Passthrough	6699
Printing *(monitored)*		
AppSocket	LZ+TFO+DRE	9100
IPP	LZ+TFO+DRE	631
SUN-Xprint	LZ+TFO+DRE	8100
Unix-Printing	LZ+TFO+DRE	515, 170
Remote-Desktop *(monitored)*		
Altiris-CarbonCopy	Passthrough	1680
Apple-NetAssistant	Passthrough	3283
Citrix-ICA	LZ+TFO+DRE	1494, 2598
ControlIT	TFO	799
Danware-NetOp	TFO	6502
Laplink-Host	TFO	1547
Laplink-PCSync	TFO	8444
Laplink-PCSync-secure	TFO	8443
MS-Terminal-Services	TFO	3389
Netopia-Timbuktu	TFO	407, 1417–1420
PCAnywhere	TFO	73, 5631, 5632, 65301
RAdmin	TFO	4899
Remote-Anything	TFO	3999, 4000
Vmware-VMConsole	TFO	902
VNC	TFO	5801–5809, 6900–6909
XWindows	TFO	6000–6063
Replication *(monitored)*		
Double-Take	LZ+TFO+DRE	1100, 1105
EMC-Celerra-Replicator	LZ+TFO+DRE	8888
MS-AD-Replication[1]	LZ+TFO+DRE+ EPM accelerator	UUID:e3514235-4b06-11d1-ab04-00c04fc2 dcd2
MS-Content-Replication-Service	TFO	560, 507

Table A-1 Predefined Traffic Policies (continued)

Classifier	WAAS Action	Destination Ports
MS-FRS[1]	LZ+TFO+DRE+ EPM accelerator	UUID:f5cc59b4-4264-101a-8c59-08002b2f 8426
Netapp-SnapMirror	LZ+TFO+DRE	10565-10569
Remote-Replication-Agent	TFO	5678
Rsync	TFO	873
SQL *(monitored)*		
Borland-Interbase	LZ+TFO+DRE	3050
IBM-DB2	LZ+TFO+DRE	523
InterSystems-Cache	LZ+TFO+DRE	1972
MS-SQL	LZ+TFO+DRE	1433
MS-SQL-RPC[1]	LZ+TFO+DRE+ EPM accelerator	UUID:3f99b900-4d87-101b-99b7-aa000400 7f07
MySQL	LZ+TFO+DRE	3306
Oracle	LZ+TFO+DRE	66, 1525, 1521
Pervasive-SQL	LZ+TFO+DRE	1583
PostgreSQL	LZ+TFO+DRE	5432
Scalable-SQL	LZ+TFO+DRE	3352
SQL-Service	LZ+TFO+DRE	156
Sybase-SQL	LZ+TFO+DRE	1498, 2638, 2439, 3968
UniSQL	LZ+TFO+DRE	1978, 1979
SSL *(monitored)*		
HTTPS	TFO	443
SSH		
SSH	TFO	22
Storage *(monitored)*		
EMC-SRDFA-IP	LZ+TFO+DRE	1748
FCIP	LZ+TFO+DRE	3225
iFCP	LZ+TFO+DRE	3420
iSCSI	LZ+TFO+DRE	3260
Streaming *(monitored)*		
Liquid-Audio	LZ+TFO+DRE	18888
MS-NetShow	LZ+TFO+DRE	1755
RTSP	LZ+TFO+DRE+ Video accelerator	554, 8554
VDOLive	LZ+TFO+DRE	7000
Systems-Management *(monitored)*		
BMC-Patrol	Passthrough	6161, 6162, 8160, 8161, 6767, 6768, 10128

Table A-1 Predefined Traffic Policies (continued)

Classifier	WAAS Action	Destination Ports
HP-OpenView	Passthrough	7426–7431, 7501, 7510
HP-Radia	LZ+TFO+DRE	3460, 3461, 3464, 3466
IBM-NetView	Passthrough	729–731
IBM-Tivoli	LZ+TFO+DRE	94, 627, 1965, 1580, 1581
LANDesk	LZ+TFO+DRE	9535, 9593–9595
NetIQ	Passthrough	2220, 2735, 10113–10116
Netopia-netOctopus	Passthrough	1917, 1921
Novell-ZenWorks	LZ+TFO+DRE	1761–1763, 517, 2544, 8039, 2037
WAAS-FlowMonitor	TFO	7878
WBEM	Passthrough	5987, 5988
Version Management *(monitored)*		
Clearcase	LZ+TFO+DRE	371
CVS	LZ+TFO+DRE	2401
VPN		
L2TP	TFO	1701
OpenVPN	TFO	1194
PPTP	TFO	1723
WAFS *(monitored)*		
CIFS	LZ+TFO+DRE+ CIFS accelerator or WAFS legacy acceleration	139, 445
Web *(monitored)*		
HTTP	LZ+TFO+DRE+ HTTP accelerator	80, 8080, 8000, 8001, 3128

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 16-58.

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

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

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
```

INDEX

A

AAA accounting
  configuring 6-33
AAA-based management systems 2-26, 6-2
acceleration
  about 1-6, 12-1
  features 1-6
  TCP adaptive buffering settings 12-43
  TCP settings 12-41
accelerators
  enabling 12-2
accounts
  creating 7-4
  creation process 7-2
  deleting 7-6
  local CLI 7-2
  roles-based 7-2
  types 7-1
  viewing 7-8
ACL
  interception 4-29
  See also IP ACL
action
  full optimization 12-36
  passthrough 12-36
  TFO only 12-36
  TFO with DRE 12-36
  TFO with LZ compression 12-36
  types 12-36
activating devices 15-31
adaptive buffering, TFO 12-43
adding
  charts 16-12
  print clusters 13-14
  printers 13-12
  print server to clients 13-24
administrative login authentication and authorization
  default 6-4
  for WAEs 6-3
  local database
    description of 6-7
  overview of 6-1
  RADIUS
    overview of 6-13
  TACACS+
    overview of 6-15
administrative login authentication failover 6-27
aggregate settings
  for print servers 13-27
alarm overload detection, enabling 9-23
alarm panel
  system dashboard window 16-3
alarms
  device reporting 16-5
alerts 16-6
aliases, for file servers 11-24
allowed protocols 13-31
application acceleration
  about 1-6, 12-1
  enabling 12-2
application classifiers
  creating 12-31
  match condition 12-34
  restoring 12-38

application definition
creating 12-30
application list, viewing 12-37
application policy
creating 12-31
creation process 12-29
position 12-40
preparation tasks 12-29
restoring defaults 12-38
types 12-33
applications
monitoring 12-39, 16-2
application traffic mix chart 16-14
assigning
application to a device 12-30
application to a device group 12-30
core cluster to file server 11-20
devices to a preposition directive 11-31
devices to device groups 3-5
devices to more than one device group 3-7
audit trail logs
viewing 6-34, 16-62
authentication
default feature values 6-4
authentication databases, types of 6-3
authentication servers
configuring 6-13, 6-15
authorization
default feature values 6-4
autodiscover 1-17
autoregistration
DHCP server requirements 2-8

B

backing up
configuration files 10-7
WAAS Central Manager 15-10
WAE devices 15-11
backup and restore
cms database 15-10
virtual blade 14-12
banners
configuring 6-10
baseline groups
configuring 3-12
creation process 3-11
customizing 3-11
switching 3-13
types 3-10
working with 3-10
BIC TCP 1-5
bootflags 15-19
browser support 2-11

C

caching, about 1-18
CDP
configuring 5-14
cdp enable command 4-40
cdp run command 4-40
Central Manager. *See* WAAS Central Manager
charts
adding 16-12
application traffic mix 16-14
customizing 16-10
descriptions 16-14
settings 16-12
CIFS 11-2
enabling 15-30
CIFS accelerator
configuring 11-9
enabling 12-2
transparent CIFS and legacy modes 11-3
CifsAO WAE Device Manager option 10-19
Cisco.com
obtaining software files from 15-3

Cisco Discovery Protocol. *See* CDP
classifier, creating 12-31
classifier report, viewing 12-38
clear ip wccp command 4-1
clear statistics all command 6-26
clear statistics authentication command 6-26
clear statistics windows-domain command 6-26
CLI user
  creating 7-4
clock
  setting 9-5
clustering in inline mode 4-50
cms database
  backup and restore procedure 15-10
cms database backup command 15-10
cms database restore command 15-11
coherency
  age-based validation 11-6
compression, about 1-5
conditions
  modifying or deleting from IP ACLs 8-6
configuration group 3-3
configuration process 11-9
congestion windows, about 5-11
connections
  creating for WAFS 11-23
  viewing TCP connections 16-45
Connections Statistics report 16-45
controlled shutdown 15-33
copy disk ftp command 15-10
core cluster
  assigning to file server 11-20
  connecting to Edge WAEs 11-23
core configuration 11-11
corrupted system images
  recovering from 15-17
CPU Statistics report 16-47
creating
  accounts 7-4
  application classifier 12-31
  application definition 12-30
  application policy 12-31
  connections for WAFS 11-23
  local user 7-4
  match condition 12-34
  new software file 15-4
  preposition directive 11-25
  preposition schedule 11-32
current software version
  determining 15-3

D

dashboard
  customizing 16-10
  device 16-9
  system 16-2
database backup 15-10
data coherency, about 11-5
data concurrency, about 11-6
data migration 2-28
data redundancy elimination, about 1-5
debug command 16-65
default status, restoring 15-12
deleting
  accounts 7-6
  device groups 3-6
  locations 3-15
  roles 7-14
  software files 15-9
  user groups 7-20
device
  alarms 16-5
  autodiscovery 1-17
  clock setting 9-5
  rebooting 15-32
Device Dashboard window 16-9

device groups
about 3-1
adding and removing devices 3-5
configuring 3-4
creating 3-3
creation process 3-2
deleting 3-6
enabling overlap 3-7
force group settings 3-8
list 3-7
overriding settings 3-8
setting configuration precedence 3-8
types 3-2
Device Home window. See Device Dashboard window
device locations
about 3-14
creating 3-14
deleting 3-15
device logs, viewing 16-63
device registration information
recovering 15-21
devices
activating 15-31
adding to device groups 3-5
adding to multiple device groups 3-7
impact of assigning to multiple groups 3-10
overriding device group settings 3-9
restarting 15-32
topology 16-44
viewing group assignments 3-6
viewing information for 16-7, 16-35, 16-45
Devices window 16-7
DHCP
configuring interfaces for 5-8
for autoregistration 2-8
interface-level 2-9
DHCP server
requirements for autoregistration 2-8
directed mode 5-16
disabling WCCP flow redirection 4-18
disconnected mode
about 11-34
configuring 11-36
requirements 11-35
disk-based software, missing
recovering from 15-20
disk encryption 15-28
disk handling
configuring error-handling methods 15-29
configuring extended object cache 15-30
disks
monitoring 16-48
Disks report 16-48
distributing drivers 13-20
DNS, configuring 5-14
domains
about 7-14
adding entities 7-15
assigning to user accounts 7-16
assigning to user groups 7-19
creating 7-15
deleting 7-17
modifying and deleting 7-17
viewing 7-17
double-byte language support 11-16
downgrading 15-3
DRE, about 1-5
drivers
distributing 13-20
installing 13-18
repository 13-16
DSCP 11-13, 11-16, 12-34
global default 12-39
dynamic shares, creating 11-21

E

edge configuration 11-14

egress methods
configuring 4-30
email server settings for reports 9-24
enable command 6-16
enabling
core file services 11-11
edge file services 11-14
optimization and accelerators 12-2
print banners 13-31
print services 13-11
SNMP 17-10
SNMP agent 17-8
traffic statistic collection 12-30
virtual blade 14-3
WCCP flow redirection 4-18
encryption
disk 15-28
enabling secure store 9-10
entities
adding to domains 7-15
EPM accelerator
enabling 12-2
errors
disk drives 15-29
EtherChannel
configuring 5-7
Exec timeout
configuring 6-11
explicit congestion notification
about 5-11
extended object cache 15-30

F

failover, for administrative login authentication 6-27
fast offline detection
about 9-23
configuring 9-22
file locking, about 11-6
file server aliases 11-24
File Server Rename utility 10-19
file servers
registering 11-17
supported 11-8
file services 11-9
about 1-7
enabling on core 11-11
enabling on edge 11-14
features 1-8
preparing for 11-8
firewall, configuring for 5-16
flash memory
corrupted 15-17
flow monitoring
configuring 16-53
force group settings 3-8
full optimization action 12-36

G

generic GRE egress method 4-30
generic routing encapsulation. *See* GRE encapsulation
Gigabit Ethernet interfaces
modifying 5-5
GRE encapsulation 4-15, 4-17
GRE packet forwarding 4-17
GRE tunnel, configuring on router 4-32
groups. *See* user groups

H

hardware clock 9-5
HTTP accelerator
configuring 12-5
enabling 12-2

I

increased buffering 1-5
inline mode 4-42
  configuring IP address 4-47
  interface settings 4-44
  serial clustering 4-50
  VLAN configuration 4-49
  VLAN ID check 4-47
inline network adapter card 4-42
intelligent message prediction 1-6
interception ACL 4-29
interface-level DHCP
  description 2-10
  note 2-8
interfaces
  manually configuring for DHCP 5-8
IP access control lists. *See* IP ACL
IP ACL
  adding conditions to 8-3
  applying to interface 8-6
  associating with application 8-6
  conditions, modifying or deleting 8-6
  configuration constraints 8-2
  creating new 8-3
  deleting 8-7
  on routers 2-24
  on WAEs 2-24
  overview 8-1
IP addresses
  multiple, configuring on single interface 5-5
  static 2-10
IP routes
  configuring 5-13
ip wccp command 4-11
ip wccp redirect-list command 4-11
ip web-cache redirect command 4-1, 4-11

K

kernel debugger
  enabling 16-63

L

Layer 2 redirection 4-17
LDAP server signing 10-11, 10-13
  configuring on a Microsoft server 6-24
  configuring on a WAE 6-25
  disabling on a WAE 6-26
  overview of 6-24
legacy mode, CIFS 11-3
licenses 9-3
line console carrier detection
  configuring 6-12
load balancing 1-20, 4-13, 5-9
local CLI accounts, about 7-2
local user, creating 7-4
locations
  about 3-14
  creating 3-14
  deleting 3-15
location tree
  viewing 3-15
logging
  configuring system logging 16-55
  message priority levels 16-57
  transaction log format B-1
  transaction logging 16-58
  viewing audit trail log 16-62
  viewing device logs 16-63
  viewing system messages 16-61
login
  WAE Device Manager 10-2
login access
  controlling 6-7

login authentication
about 2-26, 6-1
logs
severity levels in the WAE Device Manager 10-33
viewing in the WAE Device Manager 10-32
lost administrator passwords
recovering 15-19
LZ compression, about 1-5

M

management IP address 9-2
MAPI accelerator
configuring 12-8
enabling 12-2
match condition, creating 12-34
maximum segment size 12-42
message logs
viewing 16-61
message of the day settings
configuring 6-10
MIBs
supported 17-4
MIB traps
configuring using the WAE Device Manager 10-9
migration, data 2-28
missing disk-based software
recovering from 15-20
modifying configuration 13-29
monitoring
applications 12-39, 16-2
chart descriptions 16-14
chart settings 16-12
CPU utilization 16-47
creating custom reports 16-49
disk information 16-48
flows with NetQoS 16-53
predefined reports 16-35
system status 16-6
using the WAE Device Manager 10-24
with SNMP 17-1
multiple IP addresses
configuring on single interfaces 5-5

N

namespace support 1-21
NAS appliances 1-18
NAT address 9-2
NAT configuration 9-2
NetBIOS 9-2, 11-15
NetQoS monitoring 16-53
network
viewing information for 16-2
Network Time Protocol. *See* NTP
network traffic analyzer tool 16-65
NFS accelerator
enabling 12-2
notification settings
for alerts 10-15
for reports 9-24
NTP, configuring 9-5

O

obtaining software files 15-3
operation prediction and batching 1-6
optimization, enabling 12-2

P

packet forwarding method 4-15
Layer 2 redirection 4-17
Layer 3 GRE 4-17
packet return 4-16
passthrough action 12-36

passwords
changing account 7-7, 7-8
recovering administrator 15-19
PBR, about 1-20
policy-based routing
about 1-20
configuration of 4-33
overview of 2-20
verifying next-hop availability 4-39
policy report, viewing 12-38
port channel interfaces
configuring 5-7
load balancing 5-9
ports
139 2-6
bypassing 2-7
disabling 11-15
enabling 11-15
4050 2-6
445 2-6
disabling 11-15
enabling 11-15
50139 2-6
used in WAFS 2-6
position, application policy 12-40
power failure 15-17
preposition
about 11-4
checking status of 11-33
creating directive 11-25
scheduling 11-32
viewing in the WAE Device Manager 10-20
print accelerator 1-8
print banners, enabling and disabling 13-31
print clusters
about 13-4
adding 13-14
print drivers
distributing 13-20
installing 13-18
support issues 13-3
verifying distribution 13-21
printer 13-29
adding to a WAAS print server 13-12
printer misconfiguration, PostScript error 13-32
printing a test page 13-30
print servers
adding to clients 13-24
details, viewing 13-25
print services
about 1-8, 13-1
configuration process 13-7
enabling 13-11
planning 13-5
Print Services Administration GUI, about 13-28
priority levels 16-57

R

RADIUS
authentication overview 6-13
configuring server 6-13
database 6-3
default configuration 6-4
RAID 1-21
RCP services, enabling 9-4
rebooting devices 15-32
receive buffer size 12-42
recovering
device registration information 15-21
from missing disk-based software 15-20
lost administrator passwords 15-19
system software 15-17
redirection methods 4-1
registering
file servers 11-17
WAEs in the WAE Device Manager 10-6

remote login
controlling access 6-7
reports
configuring email server settings 9-24
Connection Statistics 16-45
CPU utilization 16-47
creating custom 16-49
customizing 16-10
editing 16-50
managing 16-48
predefined 16-35
scheduling 16-51
Topology 16-44
viewing custom 16-50
repository, driver 13-16
request redirection methods 4-1
rescue system image 15-17
restarting devices 15-32
restoring
application classifiers 12-38
application policies 12-38
configuration files 10-7
WAAS Central Manager 15-10
WAE devices 15-11
WAE to default condition 15-12
retransmit time multiplier
about 5-11
roles
about 7-10
assigning to user accounts 7-13
assigning to user groups 7-19
creating and managing 7-11
deleting 7-14
modifying and deleting 7-13
read-only access to services 7-11
viewing 7-14
viewing settings 7-14
roles-based accounts
about 7-2, 7-3
router
configuring WCCP transparent redirection on 4-7
router lists, defining for WCCP services 4-26

S

SACK, about 1-5
scheduling
preposition 11-32
reports 16-51
secure shell
configuring 6-8
host keys 6-9
secure store
changing key and password 9-14
configuring 9-10
disabling 9-16
enabling on Central Manager 9-12
enabling on standby Central Manager 9-13
enabling on WAE 9-13
security
disk encryption 15-28
enabling secure store 9-10
selective acknowledgement 1-5
send buffer size 12-42
send TCP keepalive 12-42
serial clustering in inline mode 4-50
service password
configuring 4-11
set ip next-hop verify-availability command 4-41
shadow copy for shared folders 11-7
show cdp neighbors command 4-39
show command utility
for troubleshooting 16-65
show version command 15-18
shutting down WCCP 4-27
Simple Network Management Protocol. *See* SNMP
site and network planning 2-4

SNMP 1-22
asset tag setting 17-20
community settings 17-15
configuration process 17-8
configuring using the WAE Device Manager 10-8
contact settings 17-20
defining custom traps 17-12
enabling 17-10
enabling SNMP agent 17-8
enabling traps 17-10
group settings 17-17
host settings 17-14
manager
creating 17-3
monitoring with 17-1
preparation 17-9
security models and security levels 17-4
supported MIBs 17-4
trap source settings 17-21
user settings 17-18
versions supported 17-3
view settings 17-16
software
recovering 15-17
software clock 9-5
software files
obtaining from Cisco.com 15-3
software licenses 9-3
software recovery
using the CD-ROM 15-12
software upgrades 15-4
for multiple devices 15-8
process 15-1
software version
determining 15-3
spooling space, default 13-15
SSL
configuring 12-11
standby Central Manager
switching to primary 15-26
standby groups
of interfaces 5-2
standby interfaces
configuring 5-2
priority settings 5-3
starting WAE components 10-5
static IP addresses 2-10
static IP routes
configuring 5-13
statistics, collecting 12-30
stopping WAE components 10-5
switching baseline group 3-13
system configuration settings 9-17
system dashboard
viewing system-wide information 16-2
system event logging
configuring 16-55
message priority levels 16-57
viewing log 16-61
system image
recovering 15-17
system message log
using 16-55
viewing 16-61
system software
recovering 15-17
system status
monitoring 16-6

T

TACACS+
authentication and authorization, overview of 6-15
database 6-3
default configuration 6-4
enable password attribute 6-16

TACACS+ server
configuring 6-15
taskbar icons 1-12
TCP
congestion windows 5-11
explicit congestion notification 5-11
parameter settings 5-9
retransmit timer 5-11
slow start 5-12
viewing connections 16-45
tcpdump command 16-65
TCP initial window size, about 1-4
TCP promiscuous mode service
overview of 2-23
Telnet services
enabling 6-10
test command for troubleshooting 16-65
tethereal command 16-65
TFO
about 1-4
TFO adaptive buffering 12-43
TFO and LZ compression action 12-36
TFO features 1-4
BIC TCP 1-5
compression 1-5
increased buffering 1-5
selective acknowledgement 1-5
TCP initial window size maximization 1-4
Windows scaling 1-4
TFO only action 12-36
TFO with DRE action 12-36
time zones
location abbreviations 9-7
parameter settings for 9-5
Topology report 16-44
track command 4-41
traffic statistics collection, enabling 12-30
traffic statistics report 16-2
chart descriptions 16-14
transaction logging 16-58
configuring 16-58
log format B-1
transparent redirection, configuring on a router 4-7
traps
defining SNMP 17-12
enabling 17-10
triggers
defining SNMP 17-12
troubleshooting
CLI commands 16-65
using show command utility 16-65
with Central Manager diagnostic tests 16-64
with TCPdump 16-65
with Tethereal 16-65
Troubleshooting Devices window 16-6
types of application policies 12-33

U

Unicode support 2-11
upgrading
device groups 15-8
process 15-1
WAAS Central Manager device 15-6
user accounts
adding domain entities 7-15
assigning to domains 7-16
audit trail logs
viewing 6-34, 16-62
changing passwords 7-7, 7-8
creating 7-4
creation process 7-2
deleting 7-6
deleting domains 7-17
domains 7-15
managing 7-8
modifying and deleting 7-6
roles

assigning to 7-13
creating 7-11
modifying and deleting 7-13
viewing 7-14
viewing 7-8
viewing domains 7-17
user authentication. *See* login authentication
user groups
about 7-18
assigning roles to 7-19
assigning to domains 7-19
creating 7-18
deleting 7-20
viewing 7-20
UTC offsets 9-8
See also GMT offsets

V

version of software 15-3
video accelerator
configuring 12-9
enabling 12-2
viewing
application list 12-37
classifier report 12-38
logs in the WAE device manager 10-32
policy report 12-38
print server details 13-25
role settings 7-14
virtual blade
backing up and restoring 14-12
configuring 14-1, 14-4
copying disk image to 14-10
enabling 14-3
starting and stopping 14-9
virtualization. See virtual blade
VLAN ID check 4-47
VLAN support 4-43

W

WAAS
benefits 1-17
interfaces 1-9
WAAS Central Manager
backing up 15-10
configuring as driver repository 13-16
restoring 15-10
upgrading 15-6
WAAS Central Manager GUI
about 1-9
accessing 1-9
components 1-10
taskbar icons 1-12
WAAS CLI, about 1-16
WAAS interfaces
CLI 1-16
Print Services Administration GUI 1-16
WAAS Central Manager GUI 1-9
WAE Device Manager GUI 1-15
WAAS networks
and IOP interoperability 2-11
network planning for 2-1
traffic redirection methods 2-18
WAAS Print Services Administration GUI, about 1-16
WAAS services, about 1-4
WAE Device Manager
about 1-15, 10-1
Configuration option 10-8
Control option for the WAE 10-4
logging out 10-3
Notifier tab 10-15
quick tour 10-2
Utilities option 10-17
workflow 10-3
WAE devices
backing up 15-11
controlled shutdown 15-33

modifying configuration properties 9-1
restoring 15-11
supported 2-11
WAE packet return 4-16
WAFS
about 1-18
configuration process 11-9
preparing for 11-8
WAFS Cache Cleanup utility 10-18
WAFS Core cluster 3-3
WAN failures, preparing for 11-34
WAN utilization setting 11-25
WCCP
about 1-19
Cisco Express Forwarding (CEF) 4-17
defining router lists 4-26
flow redirection
enabling and disabling 4-18
GRE packet return 4-30
ports used 2-6
request interception method 4-3
shutting down 4-27
WCCP-based routing
advanced configuration
for a router 4-7
configuration guidelines 4-5
overview of 2-19
wccp command 4-5
WCCP service masks
deleting 4-21
modifying 4-21, 4-23
WCCP services
displaying list of 4-5
web application filter
configuring 9-19
web browser support 2-11
Windows Authentication
checking the status in the WAE Device Manager 10-13
configuring using the WAE Device Manager 10-10
Windows name services 5-15
Windows print accelerator, about 1-8
Windows scaling, about 1-4