



Deploying Cisco Wide Area Application Services

Second Edition

Design and deploy Cisco WAN optimization and application acceleration solutions for the enterprise WAN

Deploying Cisco Wide Area Application Services, Second Edition

Joel Christner, Zach Seils, Nancy Jin

Copyright© 2010 Cisco Systems, Inc.

Published by:

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America

First Printing January 2010

Library of Congress Cataloging-in-Publication data is on file.

ISBN-13: 978-1-58705-912-4

ISBN-10: 1-58705-912-6

Warning and Disclaimer

This book is designed to provide information about deploying Cisco Wide Area Application Services (WAAS). Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Corporate and Government Sales

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact: **U.S. Corporate and Government Sales** 1-800-382-3419 corpsales@pearsontechgroup.com

For sales outside the United States please contact: **International Sales** international@pearsoned.com

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Publisher: Paul Boger

Cisco Representative: Erik Ullanderson

Associate Publisher: Dave Dusthimer

Cisco Press Program Manager: Anand Sundaram

Executive Editor: Mary Beth Ray

Copy Editor/Proofreader: Deadline Driven Publishing

Managing Editor: Patrick Kanouse

Technical Editors: Jim French, Jeevan Sharma

Senior Development Editor: Christopher Cleveland

Indexer: Angie Bess

Project Editor: Ginny Bess Munroe

Editorial Assistant: Vanessa Evans

Cover Designer: Sandra Schroeder

Book Designer: Louisa Adair

Composition: Mark Shirar



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Arionet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNF, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (0812R)

Foreword

I am pleased to write the foreword to the second edition of *Deploying Cisco Wide Area Application Services (WAAS)*. Over the past few years, WAN Optimization technology has become a standard component of enterprise networks. The benefits accruing from the use of the technology for server consolidation, simplified IT management, and improvement of the efficiency of information sharing and network utilization have earned it a place at the top of customers' buying priorities.

At Cisco, we have made several innovations to our award-winning WAAS solution that continues to expand the benefits it offers our customers. These include the use of virtualization technology—that is, Virtual Blades (VB)—to rapidly deploy a network service “anytime, anywhere,” and a variety of application specific acceleration techniques that we developed in collaboration with the leading application vendors.

At Cisco, we believe that WAN optimization technology needs to be closely integrated with the routing/VPN architecture of the enterprise network so that customers can benefit from a single, optimized, shared network fabric that delivers all applications: voice, video, and data.

The authors combine experience from their work with thousands of customers who have deployed large installations of WAAS with a deep knowledge of enterprise and service provider network design, IOS, application-aware networking technologies, and WAAS to provide a comprehensive set of best practices for customer success. I strongly recommend customers who are interested in WAN optimization and particularly Cisco WAAS to read this volume. It will help you accelerate your understanding of the solution and the benefits you can accrue.

George Kurian
Vice President and General Manager, Application Networking and Switching
Cisco Systems, Inc.

Introduction

IT organizations are realizing the benefits of infrastructure consolidation and virtualization—cost savings, operational savings, better posture toward disaster recovery—and the challenges associated. Consolidating infrastructure increases the distance between the remote office worker and the tools they need to ensure productivity—applications, servers, content, and more. Application acceleration and WAN optimization solutions such as Cisco Wide Area Application Services (WAAS) bridge the divide between consolidation and performance to enable a high-performance consolidated infrastructure.

This book is the second edition of *Deploying Cisco Wide Area Application Services*, and updates the content to reflect the innovations that have been introduced in version 4.1.3 of the Cisco Wide Area Application Services (WAAS) solution, whereas the first edition was written to version 4.0.13. Along with coverage of the key components of the Cisco WAAS solution, this edition expands on the concepts introduced in the first edition to provide a more complete understanding of the solution's capabilities, how to use them effectively, and how to manage them. This edition expands upon the first edition to include coverage for new solution components including application-specific acceleration techniques, hardware form factors, virtualization, application performance management (APM), monitoring and reporting enhancements, and workflow enhancements. Additional technical reference material is provided in the appendices to help familiarize users of version 4.0 with changes that have occurred in the command-line interface (CLI) with the introduction of the 4.1 release. A quickstart guide is provided to help users quickly deploy in a lab or production pilot environment in order to quantify the benefits of the solution. A troubleshooting guide can also be found at the end which helps associate difficulties encountered with potential steps for problem resolution.

Goals and Methods

The goal of this book is to familiarize you with the concepts and fundamentals of sizing and deploying Cisco WAAS in your environment. The book provides a technical introduction to the product, followed by deployment sizing guidelines, through integration techniques, and configuration of major components and subsystems. The intent of the book is to provide you with the knowledge that you need to ensure a successful deployment of Cisco WAAS in your environment, including configuration tips, pointers, and notes that will guide you through the process.

Who Should Read This Book?

This book is written for anyone who is responsible for the design and deployment of Cisco WAAS in their network environment. The text assumes the reader has a basic knowledge of data networking, specifically TCP/IP and basic routing and switching technologies.

As the WAAS technology continues to evolve, the content in this book will provide a solid framework to build on. Mastering the topics in this book will ensure that you can approach any WAAS design project with confidence.

How This Book Is Organized

Although this book could be read cover to cover, it is designed to be flexible and allow you to easily move between chapters and sections of chapters to cover just the material that you need to work with. Although each of the chapters builds upon the foundation laid by previous chapters, enough background information is provided in each chapter to allow it to be a standalone reference work in and of itself. Chapter 1 provides a technical examination of the Cisco WAAS product and its core capabilities, along with use cases and the “why you care” about each of the solution components. Chapters 2 through 10 are the core chapters and, although they can be covered in any order, it is recommended that they be covered sequentially for continuity. Chapter 11 provides a series of use cases for the Cisco WAAS product family, which can also provide insight into how other customers use this technology to meet their business infrastructure requirements. Appendices are provided to help augment and also summarize what is discussed in the core chapters. Following is a description of each chapter:

- **Chapter 1, “Introduction to Cisco Wide Area Application Services (WAAS):”** This chapter provides a technical examination and overview of Cisco WAAS and its core components.
- **Chapter 2, “Cisco WAAS Architecture, Hardware, and Sizing:”** This chapter discusses the Cisco WAAS appliance and router-integrated network module hardware family, positioning of each of the platforms, and system specifications that impact the design of a solution relative to the performance and scalability of each component.
- **Chapter 3, “Planning, Discovery, and Analysis:”** Planning is a critical part to any successful WAAS deployment. Spending ample time at the beginning of the project to understand the requirements, including those imposed by the existing network environment, is critical for a successful deployment. Chapter 3 gives you a head start by outlining the key topic areas that should be taken into consideration as you are planning your WAAS deployment.
- **Chapter 4, “Network Integration and Interception:”** This chapter provides an in-depth review of the network integration and interception capabilities of Cisco WAAS. The topics discussed in Chapter 4 form the foundation for the design discussions in subsequent chapters.
- **Chapter 5, “Branch Office Network Integration:”** This chapter provides a detailed discussion of the different design options for deploying Cisco WAAS in the branch office environment. Several design options are discussed, including detailed configuration examples.
- **Chapter 6, “Data Center Network Integration:”** This chapter examines the key design considerations for deploying WAAS in the data center. Sample design models and configuration examples are provided throughout the chapter. Best practices recommendations for scaling to support hundreds or thousands of remote sites are also included.

- **Chapter 7, “System and Device Management:”** This chapter walks you through the initial deployment of the Central Manager and each of the accelerator WAAS devices, including the setup script, registration, federated management, and use of management techniques such as device groups. This chapter also provides a detailed understanding of integration with centralized authentication and authorization, alarm management, an introduction to the monitoring and reporting facilities of the CM, CM database maintenance (including backup and recovery), and the XML-API.
- **Chapter 8, “Configuring WAN Optimization:”** This chapter guides you through the WAN optimization framework provided by Cisco WAAS, including each of the optimization techniques and the Application Traffic Policy manager. This chapter also examines the configuration of optimization policies, verification that policies are applied correctly, and an examination of statistics and reports.
- **Chapter 9, “Configuring Application Acceleration:”** This chapter focuses on the application acceleration components of Cisco WAAS, including configuration, verification, and how the components interact. This chapter also looks closely at how these components leverage the underlying WAN optimization framework, how they are managed, and an examination of statistics and reports.
- **Chapter 10, “Branch Office Virtualization:”** This chapter examines the virtualization capabilities provided by certain Cisco WAAS appliance devices, including configuration, management, and monitoring.
- **Chapter 11, “Case Studies:”** This chapter brings together various topics discussed in the previous chapters through several case studies. The case studies presented focus on real-world deployment examples, a discussion of the key design considerations, options, and final device-level configurations.
- **Appendix A, “WAAS Quickstart Guide:”** Appendix A provides a quickstart guide to help you quickly deploy WAAS in a proof-of-concept lab or production pilot.
- **Appendix B, “Troubleshooting Guide:”** Appendix B provides a troubleshooting guide, which helps you isolate and correct commonly encountered issues.
- **Appendix C, “4.0/4.1 CLI Mapping:”** Appendix C provides a CLI mapping quick reference to help identify CLI commands that have changed between the 4.0 and 4.1 versions.

Cisco WAAS Architecture, Hardware, and Sizing

Chapter 1, “Introduction to Cisco Wide Area Application Services (WAAS),” introduced the performance challenges created by the wide-area network (WAN) and how they are addressed by the Cisco WAAS solution. Cisco WAAS is a software component that is resident on a hardware device deployed at each location with users and servers. This hardware device, which can be deployed as a router-integrated network module for the Integrated Services Router (ISR) or as an appliance, is named either Cisco Wide-Area Application Engine (WAE) or Cisco Wide-Area Virtualization Engine (WAVE). The distinction between the two is that a WAVE device, available only as an appliance, can also provide branch office virtualization services in conjunction with WAN optimization and application acceleration. WAE devices provide only WAN optimization and application acceleration and do not provide virtualization.

This chapter provides an introduction to the Cisco WAAS hardware family, along with an in-depth examination of the hardware and software architecture. This chapter also looks at the licensing options for Cisco WAAS, positioning for each of the hardware platforms, and performance and scalability metrics for each of the platforms.

Cisco WAAS Product Architecture

The Cisco WAAS product family consists of a series of appliances and router-integrated network modules that are based on an Intel x86 hardware architecture. The product family scales from 512 MB of memory to 24 GB of memory, utilizing single-processor subsystems up to dual quad-core processor subsystems. Each Cisco WAAS device, regardless of form factor, is configured with some amount of hard disk storage and a compact flash card. The compact flash card is used for boot-time operation and configuration files, whereas the hard disk storage is used for optimization data (including object cache and Data Redundancy Elimination [DRE]), swap space, software image storage repository, and guest operating system storage in the case of WAVE devices. Having a compact flash card enables the device to remain accessible on the network should the device suffer hard drive subsystem failure for troubleshooting and diagnostics purposes

(in such a scenario, optimization and virtualization services would not be operational). Also, by using the compact flash card in this way, a WAAS device can successfully boot and become accessible on the network if no disks are available to the device.

The foundational layer of the Cisco WAAS software is the underlying Cisco Linux platform. The Cisco Linux platform is hardened to ensure that rogue services are not installed and secured such that third-party software or other changes cannot be made. The Cisco Linux platform hosts a command-line interface (CLI) shell similar to that of Cisco IOS Software, which, along with the Central Manager and other interfaces, form the primary means of configuring, managing, and troubleshooting a device or system. All relevant configuration, management, monitoring, and troubleshooting subsystems are made accessible directly through this CLI as opposed to exposing the Linux shell.

The Cisco Linux platform hosts a variety of services for WAAS run-time operation. These include disk encryption, Central Management Subsystem (CMS), interface manager, reporting facilities, network interception and bypass, application traffic policy (ATP) engine, and kernel-integrated virtualization services, as shown in Figure 2-1.

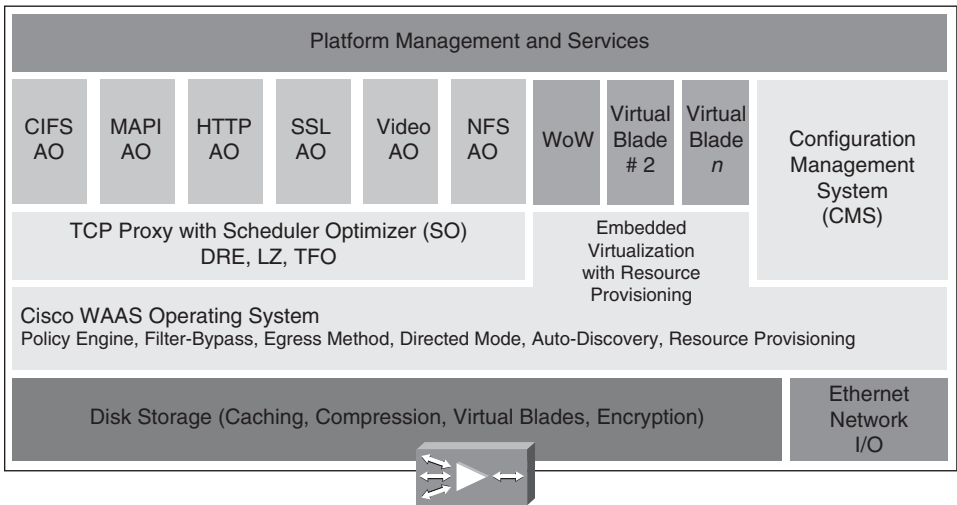


Figure 2-1 Cisco WAAS Hardware and Software Architecture

The following sections examine each of the Cisco WAAS architecture items. Cisco WAAS optimization components, including Data Redundancy Elimination (DRE), Persistent LZ Compression (PLZ), Transport Flow Optimization (TFO), and application accelerators, are discussed in detail in Chapter 1, and thus are not discussed in this chapter.

Disk Encryption

Cisco WAAS devices can be configured to encrypt the data, swap, and spool partitions on the hard disk drives using encryption keys that are stored on and retrieved from the Central Manager. The disk encryption feature uses AES-256 encryption, the strongest

commercially available encryption, and keys are stored only in the WAAS device memory after they have been retrieved from the Central Manager during the device boot process. Should a WAAS device be physically compromised or a disk stolen, power is removed from the device, which destroys the copy of the key in memory (memory is not persistent). When the hard disks are encrypted, loss of the key renders data on the disk unusable and scrambled. Keys are stored in the Central Manager database (which can be encrypted) and synchronized among all Central Manager devices for high availability. If a WAAS device is not able to retrieve its key from the Central Manager during boot time, it remains in pass-through mode until connectivity is restored or disk encryption is administratively bypassed. Additionally, the fetching of the key from the Central Manager is done over the Secure Sockets Layer (SSL)-encrypted session that is used for message exchanges between the WAAS devices and the Central Manager devices.

Central Management Subsystem

CMS is a process that runs on each WAAS device, including accelerators and Central Managers. This process manages the configuration and monitoring components of a WAAS device and ensures that each WAAS device is synchronized with the Central Manager based on a scheduler known as the Local Central Manager (LCM) cycle. The LCM cycle is responsible for synchronizing the Central Manager CMS process with the remote WAAS device CMS process to exchange configuration data, fetch health and status information, and gather monitoring and reporting data. The CMS process is tied to a management interface configured on the WAAS device known as the primary interface, which is configured on the WAAS device CLI prior to registration to the Central Manager. Any communication that occurs between WAAS devices for CMS purposes is done using SSL-encrypted connections for security.

Interface Manager

The Cisco WAAS device interface manager manages the physical and logical interfaces that are available on the WAAS device. Each WAAS device includes two integrated Gigabit Ethernet interfaces (including the network modules, one interface is internal and shares connectivity to a peer interface in the router through the router backplane, the other is external and can be cabled to a LAN switch, similar to an appliance). Each WAAS appliance has expansion slots to support one or more additional feature cards, such as the inline bypass adapter, which has two two-port fail-to-wire pairs. The interface manager also provides management over logical interfaces that can be configured over physical interfaces. Logical interfaces include active/standby interfaces, where one physical interface is used as a primary interface and a second interface is used as a backup in the event the primary interface fails. Another logical interface is the PortChannel interface, which can be used to team WAAS device interfaces together for the purposes of high availability and load balancing. It should be noted that active/standby interfaces are used when WAAS device interfaces connect to separate switches, whereas PortChannel interfaces are used when the WAAS device interfaces connect to the same switch.

Monitoring Facilities and Alarms

Cisco Linux provides an interface for the Cisco WAAS software to use for purposes of monitoring and generating alarms. Cisco WAAS supports the Simple Network Management Protocol (SNMP) versions 1, 2c, and 3, and a host of Management Information Bases (MIB) that provide complete coverage over the health of each individual WAAS device. Cisco WAAS also supports the definition of up to four syslog servers, which can be used as alarm recipients when syslog messages are generated. The WAAS Central Manager also has an alarm dashboard, which is described in Chapter 7, “System and Device Management.” The Central Manager makes an application programming interface (API) available for third-party visibility systems, which is also discussed in Chapter 7, Chapter 8, “Configuring WAN Optimization,” and Chapter 9, “Configuring Application Acceleration.” Transaction logs can be configured to be stored on each of the accelerator devices in the network for persistent retention of connection statistics, which might be useful for troubleshooting, debugging, or analytics purposes. Transaction logs are not covered in this book, but a full reference on their usage can be found in the Cisco WAAS documentation.

Note The alarm book (which covers syslog messages, SNMP traps, and Central Manager dashboard alarms), error book (which covers console messages), and product documentation can be downloaded from Cisco.com at <http://www.cisco.com/cgi-bin/tablebuild.pl/waas41>.

Network Interception and Bypass Manager

The network interception and bypass manager is used by the Cisco WAAS device to establish relationships with intercepting devices where necessary and ensure low-latency bypass of traffic that the WAAS device is not intended to handle. The Web Cache Coordination Protocol version 2 (WCCPv2) is a protocol managed by the network interception and bypass manager to allow the WAAS device to successfully join a WCCPv2 service group with one or more adjacent routers, switches, or other WCCPv2-capable server devices. WCCPv2 is discussed in more detail in Chapter 4, “Network Integration and Interception.” Other network interception options, which are also discussed in Chapter 4, include policy-based routing (PBR), physical inline interception, and Application Control Engine (ACE). As flows are intercepted by the WAAS device and determined to be candidates for optimization, those flows are handed to the Application Traffic Policy (ATP) engine to identify what level of optimization and acceleration should be applied based on the configured policies and classifier matches. The ATP is discussed in the next section, and Chapter 8 and Chapter 9 discuss the configuration and management of policies.

Application Traffic Policy Engine

Although the foundational platform component of Cisco WAAS is Cisco Linux, the foundational optimization layer of the Cisco WAAS software (which is as much a component of the Cisco Linux platform as it is the software) is the ATP engine. The ATP is responsible for examining details of each incoming flow (after being handled by the interception and bypass mechanisms) in an attempt to identify the application or protocol associated with the flow. This association is done by comparing the packet headers from each flow against a set of predefined, administratively configured, or dynamic classifiers, each with its own set of one or more match conditions. Flows that do not have a match with an existing classifier are considered “other” traffic and are handled according to the policy defined for other traffic, which indicates that there are no classifier matches and that the default policy should be used.

When a classifier match is found, the ATP examines the policy configuration for that classifier to determine how to optimize the flow. The ATP also notes the application group to which the classifier belongs to route statistics gathered to the appropriate application group for proper charting (visualization) and reporting. The configured policy dictates which optimization and acceleration components are enacted upon the flow and how the packets within the flow are handled. The list of configurable elements within a policy include the following:

- **Type of policy:** Defines whether the policy is a basic policy (optimize, accelerate, and apply a marking), Wide Area File Services Software (WAFS) transport (used for legacy mode compatibility with WAAS version 4.0 devices), and end-point mapper (EPM, used to identify universally-unique identifiers for classification and policy).
- **Application:** Defines which application group the statistics should be collected into, including byte counts, compression ratios, and others, which are then accessible via the WAAS device CLI or Central Manager.
- **Action:** Defines the WAN optimization policy that should be applied to flows that match the classifier match conditions. This includes:
 - **Passthrough:** Take no optimization action on this flow
 - **TFO Only:** Apply only TCP optimization to this flow, but no compression or data deduplication
 - **TFO with LZ Compression:** Apply TCP optimization to this flow, in conjunction with persistent LZ compression
 - **TFO with Data Redundancy Elimination:** Apply TCP optimization to this flow, in conjunction with data deduplication
 - **Full Optimization:** Apply TCP optimization, persistent LZ compression, and data duplication to this flow
- **Accelerate:** Accelerate the traffic from within this flow using one of the available application accelerators. This provides additional performance improvement above

and beyond those provided by the WAN optimization components defined in Action and includes (the capabilities are described in detail in Chapter 1):

- **MS Port Mapper:** Identify application based on its universally unique identifier, which allows WAAS to appropriately classify certain applications that use server-assigned dynamic port numbers
- **Common Internet File System (CIFS):** Acceleration for Microsoft file-sharing environments
- **HTTP:** Acceleration for intranet and Internet applications that use the hypertext transfer protocol
- **NFS:** Acceleration for UNIX file-sharing environments
- **MAPI:** Acceleration for Microsoft Exchange e-mail, calendaring, and collaboration environments
- **Video:** Acceleration for Windows Media over RTSP streams
- **Position:** Specify the priority order of this policy. Policies are evaluated in priority order, and the first classifier and policy match determines the action taken against the flow and where the statistics for that flow are aggregated.
- **Differentiated Services Code Point (DSCP) Marking:** Apply a DSCP value to the packets in the flow. WAAS can either preserve the existing DSCP markings or apply a specific marking to the packets matching the flow based on the configuration of this setting.

Settings configured in the policy are employed in conjunction with one another. For instance, the CIFS policy is, by default, configured to leverage the CIFS accelerator prior to leveraging the “full optimization” (DRE, PLZ, TFO) capabilities of the underlying WAN optimization layer. This can be coupled with a configuration that applies a specific DSCP marking to the packets within the flow. This is defined in a single policy, thereby simplifying overall system policy management. Classifiers within the ATP can be defined based on source or destination IP addresses or ranges, TCP port numbers or ranges, or universally-unique identifiers (UUID). The ATP is consulted only during the establishment of a new connection, which is identified through the presence of the TCP synchronize (SYN) flag which occurs within the first packet of the connection. By making a comparison against the ATP using the SYN packet of the connection being established, the ATP does not need to be consulted for traffic flowing in the reverse direction, as the context of the flow is established by all WAAS devices in the path between the two endpoints and applied to all future packets associated with that particular flow. In this way, classification performed by the ATP is done once against the three-way handshake (SYN, SYN/ACK packets) and is applicable for both directions of traffic flow.

Figure 2-2 shows how the ATP engine interacts with a flow and a particular policy. For more information on ATP, including configuration, please see Chapter 8 and Chapter 9.

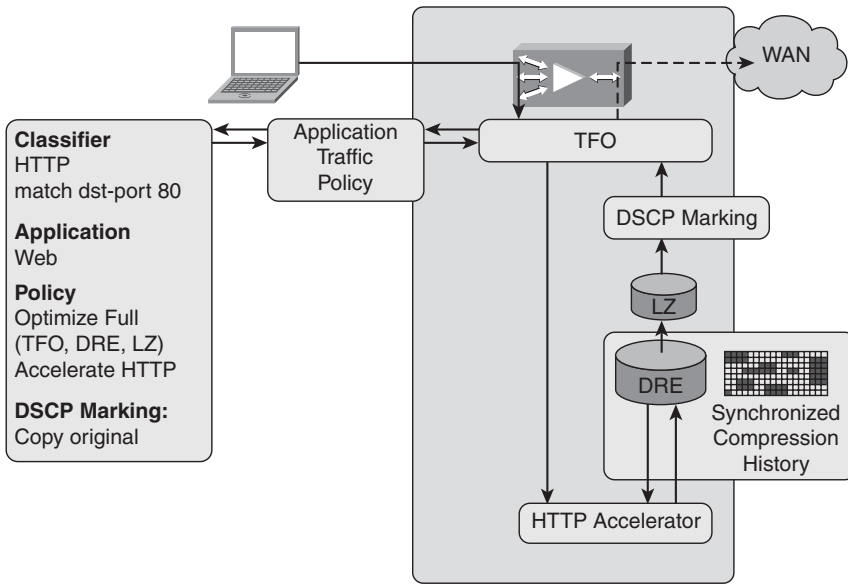


Figure 2-2 Connection Interaction with Application Traffic Policy

Virtual Blades

Cisco WAAS utilizes Kernel-based Virtual Machine (KVM) technology from Red Hat (via the Qumranet acquisition) to allow the WAVE appliance (and the WAE-674) to host third-party operating systems and applications. As of version 4.1.3, Microsoft Windows Server, versions 2003 and 2008, are supported for installation on the WAAS Virtual Blade (VB) architecture, and certain configurations can be bundled and packaged within the WAVE configuration with full support from the Cisco Technical Assistance Center (TAC). This configuration includes Microsoft Windows Server 2008 Core, Active Directory read-only domain controller, DNS server, DHCP server, and print server. The WAAS VB architecture helps enable customers to further consolidate infrastructure by minimizing the number of physical servers required in the branch office for those applications which are not good candidates for centralization into a data center location.

Hardware Family

The current Cisco WAAS hardware family consists of three router-integrated network modules, two desktop appliance models, and four rack-mounted appliance models. With such a diverse hardware portfolio, Cisco WAAS can be deployed in each location with the appropriate amount of optimization capacity for the needs of the users or servers in that particular location. This section examines the specifics of each of the current and legacy hardware platforms and positioning of each. Performance and scalability metrics for each are examined later in this chapter, along with best practices around accurately sizing a Cisco WAAS deployment.

Router-Integrated Network Modules

The Cisco WAAS router-integrated network modules are designed to provide optimization services for the remote branch office or enterprise edge. These modules, which are single-processor systems based on the Network Module Enhanced (NME) hardware, can occupy an empty or available NME-capable slot in a Cisco Integrated Services Router (ISR), including models 2811, 2821, 2851, 3825, and 3845. The ISR is an ideal platform for the branch office in that it provides a converged service platform for the remote office, including routing, switching, wireless, voice, security, and WAN optimization in a single chassis (platform, software version, and slot capacity dependent). In addition, the ISR provides a strong foundation for application performance management (APM) solutions in that along with WAAS, other performance-related features can be configured, including quality of service (QoS) for network provisioning, Performance Routing (PfR) for optimal path selection and network utilization, and NetFlow for visibility into traffic distribution, throughput, and other metrics.

Figure 2-3 shows a picture of the Cisco NME-WAE family of WAAS integrated network modules and the ISR family.



Figure 2-3 *Cisco ISR Family and WAAS Network Modules*

The Cisco NME-WAE family includes three models: the NME-WAE-302, NME-WAE-502, and NME-WAE-522. Each network module has a single hard disk with capacity ranging from 80 to 160 GB. With only a single drive, the NME-WAE is not capable of Redundant Array of Inexpensive Disks (RAID). NME-WAE devices integrate into the network using WCCPv2 as a means of interception (Policy-Based Routing [PBR] can also be used, but WCCPv2 is preferred). Both methods of integration and interception are discussed in Chapter 4. The NME-WAE family does not provide support for virtualization in the branch office; a WAVE appliance model or WAE-674 is required for virtualization support. Each NME-WAE has two network interfaces:

- **One internal:** Connected to the ISR backplane, which communicates with an internal network interface on the ISR

- **One external:** Accessible through the front of the module, which can be attached to a LAN switch

Figure 2-4 shows the architecture of the NME, internal and external interfaces, and intersection points between the NME and the ISR.

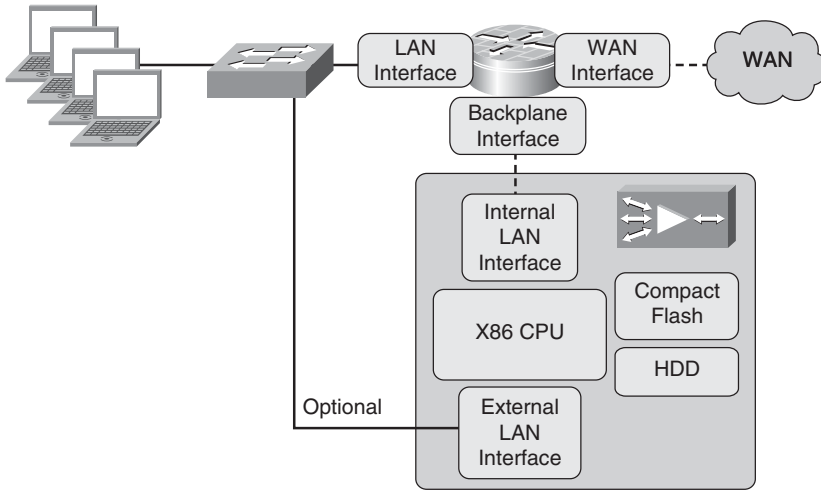


Figure 2-4 Cisco WAAS Network Module Architecture

NME-WAE Model 302

The Cisco NME-WAE model 302 (NME-WAE-302) is designed for customers who want to employ only basic WAN optimization capabilities, which are permitted through the use of the Transport license (licensing is discussed later in this chapter). These capabilities include the ATP engine, DRE, PLZ, and TFO. This module is not capable of running the advanced services enabled by the Enterprise license (discussed later in the chapter), including application layer acceleration or disk encryption. The NME-WAE-302 is a single-processor system with 512 MB of RAM and a single 80-GB hard disk.

NME-WAE Model 502

The Cisco NME-WAE model 502 (NME-WAE-502) is designed for customers who want to employ WAN optimization capabilities and application acceleration features for an enterprise edge location. The NME-WAE-502 can be configured with the Enterprise license, providing full WAN optimization functionality, application acceleration functionality, and other features enabled by the Enterprise license including disk encryption and NetQoS integration. The NME-WAE-502 is a single-processor system with 1 GB of RAM and a single 120-GB hard disk. The NME-WAE-502 is capable of supporting a larger number of users than the NME-WAE-302, as discussed in the “Performance and Scalability Metrics” section later in this chapter.

NME-WAE Model 522

The Cisco NME-WAE model 522 (NME-WAE-522) is designed for customers who want to employ appliance-equivalent functionality to an enterprise edge location in the ISR. The NME-WAE-522 supports the full suite of Enterprise license features, including all WAN optimization and application acceleration capabilities. The NME-WAE-522 is a single-processor system with 2 GB of RAM and a 160-GB hard disk, serving as the most powerful network module available as of this writing.

Appliances

The Cisco WAAS appliance family is designed to be deployed in a location of any size, including the small branch office, campus networks, or the largest of enterprise data center networks. The Cisco WAAS appliance family includes the WAE and the newer WAVE devices. Cisco WAVE appliances are current-generation and provide support for branch office virtualization, whereas WAE appliances (with the exception of the WAE-674) do not. The Cisco WAE family includes models 512, 612, 674, 7341, and 7371, and the Cisco WAVE family includes models 274, 474, and 574. WAE appliance models 512 and 674, along with WAVE appliance models 274, 474, and 574, are targeted toward branch office deployments, whereas the WAE appliance models 674, 7341, 7371 are targeted toward regional office and data center deployments. The WAE-674 is a hybrid device that is commonly used for larger branch offices (and those where virtualization is required), but works equally well as a data center device where virtualization is not used. This should not imply that the device characterization presented is fixed; devices should be placed in locations according to performance and scalability sizing and feature requirements.

The WAE appliance models 512, 612, 674, 7341, and 7371, along with WAVE appliance model 574, each have externally accessible hard disk drives and RAID support (some models support hot-swappable disk drives). WAVE appliance models 274 and 474 do not have externally accessible hard disk drives, and with a single hard disk drive, do not support RAID.

Each WAE and WAVE appliance has two built-in Gigabit Ethernet interfaces, which can be deployed independently of one another or as a pair in either an active/standby configuration or PortChannel configuration. Such interface configurations are discussed in Chapter 5, “Branch Office Network Integration,” and Chapter 6, “Data Center Network Integration.” The WAE and WAVE appliance families both have one or more Peripheral Component Interconnect (PCI) expansion slots that support installation of additional feature cards, such as the physical in-path interception card. Each WAE or WAVE appliance can be deployed using a variety of network interception techniques, including physical inline interception, WCCPv2, PBR, and ACE (all are described in Chapter 4). Any appliance model can be used as a core (data center) or edge (branch office) device, although performance and scalability recommendations presented in this chapter must be followed. Figure 2-5 shows an image of the Cisco WAE appliance family, and Figure 2-6 shows an image of the Cisco WAVE appliance family.



Figure 2-5 *Cisco WAAS WAE Appliance Family*



Figure 2-6 *Cisco WAAS WAVE Appliance Family*

Note The WAE model 7326 is end-of-life and is not covered in this section; however, its performance and scalability metrics are covered in this chapter to support those who have already deployed these devices in their networks and wish to continue using them with WAAS v4.1.

WAVE Model 274

The Cisco WAVE model 274 (WAVE-274) is a single-processor desktop model that is designed for deployment in small and medium-sized branch office locations or small data center locations. The WAVE-274 is configured with 3 GB of RAM. The WAVE-274 provides full WAN optimization and application acceleration capabilities and supports virtualization with up to two VBs. The WAVE-274 can be configured with any license available for WAAS. The WAVE-274 includes a single 250-GB SATA2 hard disk drive, and therefore does not support RAID. The WAVE-274 includes an inline card (with support for one WAN link) and the Enterprise license (discussed in the “Licensing” section of this chapter).

WAVE Model 474

The Cisco WAVE model 474 (WAVE-474) is a single-processor desktop model that is designed for deployment in small- and medium-sized branch office locations or small data center locations. Like the WAVE-274, the WAVE-474 is configured with 3 GB of

RAM. The WAVE-474 provides full WAN optimization and application acceleration capabilities and supports virtualization with up to two VBs. The WAVE-474 can be configured with any license available for WAAS. The WAVE-474 includes a single 250-GB SATA2 hard disk drive, and with a single drive, it does not support RAID. The WAVE-474 is similar to the WAVE-274, but supports a larger number of optimized TCP connections and higher levels of WAN bandwidth. The WAVE-474 includes an inline card (with support for two WAN links) and the Enterprise license (discussed in the “Licensing” section of this chapter).

WAE Model 512

The Cisco WAE model 512 (WAE-512) is a single-processor rack-mount system that is designed for deployment in small- and medium-sized branch office locations or small data center locations. The WAE-512 can be configured with 1 or 2 GB of RAM. In either configuration, the WAE-512 can provide full WAN optimization and application acceleration capabilities, but does not support virtualization. With an increase in memory configuration, the WAE-512 supports a larger number of optimized TCP connections and a greater amount of WAN bandwidth. Regardless of memory configuration, the WAE-512 can be configured with the Transport, Enterprise, or Video license. The WAE-512 supports two 250-GB SATA2 hard disk drives, which are configured automatically for software RAID-1.

WAVE Model 574

The Cisco WAVE model 574 (WAVE-574) is a quad-core rack-mount system that is designed for deployment in large branch office locations or small data center locations. The WAVE-574 can be configured with either 3 GB or 6 GB of RAM and either one or two 500 GB SATA hard disk drives. With two drives, the system is configured automatically for software RAID-1. The 6 GB RAM configuration affords the WAVE-574 support for increased WAN bandwidth and optimized TCP connections and enables the 574 to increase its VB support from two to six (assuming 512MB of RAM is allocated for each VB). The WAVE-574 supports the full breadth of features and capabilities offered by any available Cisco WAAS license.

WAE Model 612

The Cisco WAE model 612 (WAE-612) is a dual-core processor rack-mount system that is designed for deployment in medium-sized branch office locations or medium-sized data center locations. The WAE-612 can be configured with 2 GB or 4 GB of RAM (4 GB of RAM provides greater WAN bandwidth support and higher optimized TCP connection counts) and, in any configuration, supports the full breadth of features and capabilities offered by the Transport, Enterprise, and Video licenses. The WAE-612 supports two 300-GB SAS hard disk drives, which are configured automatically for software RAID-1 and are hot-swap capable.

WAE Model 674

The Cisco WAE model 674 (WAE-674) is a quad-core rack-mount system that is designed for deployment in large branch office locations or medium to large data center locations. The WAE-674 can be configured with either 4 GB or 8 GB of RAM and three 300 GB SAS hard disk drives, which are capable of hot-swap. The 4 GB RAM configuration affords the WAE-674 support for up to two VBs, and the 8 GB RAM configuration affords the WAE-674 support for up to six VBs. Additionally, the increased memory configuration provides support for a greater amount of WAN bandwidth and optimized TCP connections. The WAE-674 is unique in that it can be configured with or without VB support (the only device in the hardware family that can), and when configured without VB support, the WAE-674 can support an even higher level of WAN bandwidth and optimized TCP connections. The reason for this level of configurability is the unique position of the WAE-674, which can be used for branch offices and data centers of virtually any size. The WAE-674 supports the full breadth of features and capabilities offered by any available Cisco WAAS license.

WAE Model 7341

The Cisco WAE model 7341 (WAE-7341) is a single quad-core rack-mount system (four processors) that is designed for deployment in large enterprise data centers. The WAE-7341 includes 12 GB of RAM and four 300-GB Serial-Attached SCSI (SAS) hard disk drives, which are configured automatically for hardware RAID-5 and support hot-swap. The WAE-7341 supports the full breadth of features and capabilities offered by the Transport, Enterprise, and Video Cisco WAAS licenses, but not virtualization.

WAE Model 7371

The Cisco WAE model 7371 (WAE-7371) is a dual quad-core rack-mount system (eight processors) that is designed for deployment in the largest of enterprise data centers and under the most demanding conditions. The WAE-7371 includes 24 GB of RAM and six 300-GB SAS hard disk drives, which are configured automatically for hardware RAID-5 and support hot-swap. The WAE-7371 supports the full breadth of features and capabilities offered by the Transport, Enterprise, and Video Cisco WAAS licenses, but not virtualization.

Licensing

Each Cisco WAAS device, whether it is an appliance (WAE or WAVE) or a router-integrated network module, must be configured with one or more licenses. This license dictates what features are permitted to be configured on the device. Licenses are not enforced in WAAS; however, licenses can only be applied to platforms that support the

particular license in question. Four licenses exist for Cisco WAAS and configuration of licenses are discussed in Chapter 7:

- **Transport license:** Enables a WAAS device to apply only basic WAN optimization capabilities. It supports use of TFO, DRE, and PLZ. WAAS devices configured with the Transport license cannot provide Enterprise license features including application-acceleration capabilities, disk encryption, or any other features provided by other licenses. WAAS devices configured with the Transport license can, however, register with and be managed and monitored by a WAAS device configured as a Central Manager. The Transport license is supported by all Cisco WAAS hardware platforms.
- **Enterprise license:** Allows a WAAS device to apply all the WAN optimization provided by the Transport license and all the application acceleration functionality with the exception of Video (which is licensed separately). Additionally, the Enterprise license enables support for disk encryption and NetQoS integration. Like the Transport license, WAAS devices configured with the Enterprise license can register with and be managed and monitored by a WAAS device configured as a Central Manager. Configuration of a WAAS device as a Central Manager requires the Enterprise license. The Enterprise license is supported by all Cisco WAAS hardware platforms with the exception of the network module model 302 (NME-302).
- **Video:** Allows a WAAS device to apply stream splitting to Windows Media over Real-Time Streaming Protocol (RTSP) traffic. The Video license is commonly applied in conjunction with the Enterprise license. The Video license is supported by all Cisco WAAS hardware platforms with the exception of the network module model 302 (NME-302).
- **Virtual-Blade:** Allows a WAAS device to host third-party operating systems and applications in one or more VBs in the branch office, including Microsoft Windows Server. The Virtual-Blade license is supported on all Cisco WAVE appliances in addition to the WAE model 674.

Performance and Scalability Metrics

Design of a Cisco WAAS solution involves many factors, but the cornerstone of the solution design is based on the performance and scalability metrics required for the solution as a whole and for each individual location where WAAS is deployed. Every component in an end-to-end system has a series of static and dynamic system limits. For instance, a typical application server might be limited in terms of the number of connections it can support, disk I/O throughput, network throughput, CPU speed, or number of transactions per second. Likewise, each Cisco WAAS device has static and dynamic system limits that dictate how and when a particular WAAS device is selected for a location within an end-to-end design. This section examines the performance and scalability metrics of the Cisco WAAS hardware family, and provides a definition of what each item is and how it is relevant to a localized (per location) design and an end-to-end system design.

The static and dynamic limits referred to are used as a means of identifying which device is best suited to provide services to a particular location in the network. The

device might be deployed as an edge device, where it connects to potentially many peer devices in one or more data center locations, or as a core device, where it serves as an aggregation point for many connected edges. WAAS devices can also be deployed as devices to optimize links between data center locations, where devices on each side are realistically core devices. A fundamental understanding of the performance and scalability metrics is paramount in ensuring a sound design. Although WAAS devices have no concept of “core” or “edge,” the deployment position within the network has an effect on the type of workload handled by a device and should be considered—primarily as it relates to TCP connection count and peer fan-out (how many peers can connect to a device for the purposes of optimization). This section examines each of the performance and scalability system limits, both static and dynamic, that should be considered. These include device memory, disk capacity, the number of optimized TCP connections, WAN bandwidth and LAN throughput, the number of peers and fan-out, and the number of devices managed.

Device Memory

The amount of memory installed in a device dictates the level of performance and scalability the device can provide. As the memory capacity increases, the ability of a WAAS device to handle a larger number of connections, a larger addressable index space for compression, or a longer history of compression data also increases. Having larger amounts of memory also enables the WAAS device to run additional services, such as application acceleration, disk encryption, or virtualization, and positions the device to accept additional features that might be introduced in future software releases.

The NME-WAE family members have fixed memory capacity and cannot be upgraded. Thus, the system limits for the NME-WAE family are static. From the WAE appliance family, the 7341 and 7371 have fixed memory configurations. However, the WAE-512, WAE-612, and WAE-674 have configurable memory options, in that:

- The WAE-512 can be configured with 1 GB or 2 GB of memory.
- The WAE-612 can be configured with 2 GB or 4 GB of memory.
- The WAE-674 can be configured with 4 GB or 8 GB of memory.

For devices that support flexible memory configuration (such as the WAE-512, WAE-612, and WAE-674), higher levels of WAN bandwidth can be realized, along with an increase in the number of optimized TCP connections that can be handled concurrently by that device. For virtualization-capable platforms, a larger number of VBs can be supported. The WAVE appliance family models 274 and 474, like the network modules, are fixed configuration and do not support a memory upgrade, whereas the 574 model—like the WAE 512, 612, and 674—does support memory configuration (either 3 GB or 6 GB).

The amount of installed memory directly impacts what license is supported on each of the device models. The Transport license can be configured on any WAAS hardware model. WAAS hardware models that have 1 GB of memory or more (all do except the NME-WAE-302) can be configured with the Enterprise license, which allows the WAAS device to operate all of the Enterprise license features.

Previous versions of Cisco WAAS (version 4.0.x and version 4.1.x when using legacy mode compatibility) had distinct *core* and *edge* CIFS acceleration services. With legacy mode, a device with 1 GB of RAM can support only edge services for CIFS, whereas a device with 2 GB of RAM or more can support edge or core services, or both together. As of Cisco WAAS version 4.1.1, this deployment mode is no longer required unless interoperability with version 4.0.x is required. Generally speaking, most customers upgrade the entire network in a short and well-defined period of time and can take advantage of the simplified deployment model provided in 4.1.x, which does not have such restrictions.

Disk Capacity

Optimization services in the Cisco WAAS hardware family leverage both memory and disk. From a disk perspective, the larger the amount of available capacity, the larger the amount of optimization history that can be leveraged by the WAAS device during runtime operation. For instance, an NME-WAE-502 has 120 GB of physical disk capacity, of which 35 GB is available for use by DRE for compression history. With 35 GB of compression history, one can estimate the length of the compression history given WAN conditions, expected network utilization, and assumed redundancy levels.

Table 2-1 shows how the length of the compression history can be calculated for a particular WAAS device, along with an example. This example assumes a T1 WAN that is

Table 2-1 *Calculating Compression History*

Step	Action	Example Result
1	Convert WAN capacity to bytes (divide the number of bits per second by 8)	$(T1 = 1.544 \text{ Mbps}) / 8 = 193 \text{ KBps}$
2	Identify maximum WAN throughput for a given day (convert from seconds to minutes, to hours, to a single day)	$193 \text{ KB/sec} * 60 \text{ sec/min}$ $11.58 \text{ MB/min} * 60 \text{ min/hr}$ $694.8 \text{ MB/hr} * 24 \text{ hr/day}$ Total 16.68 GB/day
3	Identify WAN throughput given utilization (multiply by the number of hours and utilization per hour)	$(694.8 \text{ MB/hr} * 8 \text{ hours}) * 75\% \text{ utilization} = 4.168 \text{ GB}$ $(694.8 \text{ MB/hr} * 16 \text{ hours}) * 50\% \text{ utilization} = 5.56 \text{ GB}$ Total = 9.72 GB/day
4	Identify WAN throughput given utilization and expected redundancy (multiply daily throughput by expected redundancy or compressibility)	$9.72 \text{ GB/day} * .25 \text{ (as } .75 \text{ is } 75\% \text{ redundancy)} =$ 2.43 GB/day
5	Calculate compression history (divide capacity by daily throughput)	Storage capacity of unit divided by daily throughput $35 \text{ GB} / 2.43 \text{ GB/day} =$ 14.4 days of history

75 percent utilized during business hours (75 percent utilization over 8 hours per day) and 50 percent utilized during nonbusiness hours (16 hours per day), and assumes that data traversing the network is 75 percent redundant (highly compressible by DRE). This table also assumes an NME-WAE-502 with 35 GB of allocated capacity for DRE compression history.

It is generally recommended that, at minimum, five days of compression history be available in a WAAS device to better ensure that substantial performance improvements are possible. In the example in Table 2-1, the NME-WAE-502 contains enough storage capacity to provide an effective compression history of two weeks. In most cases, users tend to access data that is newer more frequently, whereas older data is accessed less frequently. Because of this, having five days worth of compression history could even be considered overkill.

The disk capacity available to a WAAS device is split among five major components:

- **DRE compression history:** This capacity is used for storing DRE chunk data and signatures.
- **CIFS cache:** This capacity is preallocated on all devices using the Enterprise license.
- **Print services:** This capacity is preallocated for print spool capacity. Print services require that the Enterprise license be configured and that CIFS edge services be configured, which implies that legacy mode is being used. In cases where print services are configured, the 1 GB of disk capacity is allocated. Given that 1 GB is a fraction of the total storage capacity of a device, it is not accounted for in Table 2-2.
- **Platform services:** This capacity is preallocated for operating system image storage, log files, and swap space.
- **Virtual Blades:** This capacity is preallocated for any guest operating systems and applications that are installed to run in a WAAS VB.

Table 2-2 shows the storage allocation for each WAAS device for each of these components.

Number of Optimized TCP Connections

Each WAAS device has a static number of TCP connections that can be optimized concurrently. Each TCP connection is allocated memory and other resources within the system, and if the concurrently optimized TCP connection static limit is met, additional connections are handled in a pass-through fashion. Adaptive buffering (memory allocation) is used to ensure that more active connections are allocated additional memory, and less active connections are only allocated the memory they require.

The TCP connection limit of each WAAS device can be roughly correlated to the number of users supported by a given WAAS device model, but note that the number of TCP connections open on a particular node can vary based on user productivity, application behavior, time of day, and other factors. It is commonly assumed that a user will have 5 to

Table 2-2 *Disk Capacity Allocation per Platform*

Platform	Total Usable Capacity	DRE	CIFS	VBs
NME-WAE-302	80 GB	30 GB	0 GB	0 GB
NME-WAE-502	120 GB	35 GB	49 GB	0 GB
NME-WAE-522	160 GB	67 GB	67 GB	0 GB
WAVE-274	250 GB	40 GB	120 GB	35 GB
WAVE-474	250 GB	60 GB	120 GB	35 GB
WAE-512-1GB	250 GB RAID-1	60 GB	120 GB	0 GB
WAE-512-2GB	250 GB RAID-1	80 GB	100 GB	0 GB
WAVE-574-3GB	500 GB RAID-1	80 GB	120 GB	60 GB
WAVE-574-6GB	500 GB RAID-1	120 GB	120 GB	180 GB
WAE-612-2GB	300 GB RAID-1	100 GB	120 GB	0 GB
WAE-612-4GB	300 GB RAID-1	120 GB	120 GB	0 GB
WAE-674-4GB	600 GB RAID-5	120 GB	120 GB	120 GB
WAE-674-8GB	600 GB RAID-5	150 GB (with VB) 320 GB (without VB)	120 GB	200 GB (with VB) 0 GB (without VB)
WAE-7326	900 GB RAID-1	320 GB	230 GB	0 GB
WAE-7341	900 GB RAID-5	500 GB	230 GB	0 GB
WAE-7371	1500 GB RAID-5	1 TB	230 GB	0 GB

15 connections open at any given time, with roughly 6 to 10 of those connections requiring optimization. If necessary, policies can be adjusted on the WAAS Central Manager to pass through certain applications that might realize only a small amount of benefit from WAAS. This type of change could potentially help increase the number of users that can be supported by a particular WAAS device.

Table 2-3 shows the optimized TCP connection capacity per device model.

Table 2-3 *Optimized TCP Connection Capacity per Platform*

Network Module	Connection Capacity	Appliance	Connection Capacity
NME-WAE-302	250	WAVE-274	200
NME-WAE-502	500	WAVE-474	400
NME-WAE-522	800	WAE-512-1GB	600
		WAE-512-2GB	1200
		WAVE-574-3GB	750
		WAVE-574-6GB	1300
		WAE-612-2GB	1600
		WAE-612-4GB	4800
		WAE-674-4GB	2000
		WAE-674-8GB (with VB)	4000
		WAE-674-8GB (without VB)	6000
		WAE-7326	5000
		WAE-7341	12,000
		WAE-7371	50,000

The number of connections a typical user has in a location can be determined by using tools that exist in the operating system of the user's workstation. Although the estimate of six to ten optimized TCP connections is accurate for the broad majority of customers, those that wish to more accurately determine exactly how many connections a typical user has open at any given time can do so.

Microsoft provides two methods for determining the number of connections that are open on a given computer. The first is through the Command Prompt program **netstat**. By opening a Command Prompt window (click **Start > Run**, then type **cmd** and click **Ok**) and typing the command **netstat**, you can see a list of the open connections from the computer to all of the other endpoints to which that computer is connected. Notice the connections that are in the state of **ESTABLISHED**. These connections are currently open and in use and have not yet been closed. In many cases, the protocol associated with the connection is listed next to the foreign address, but some might not be. From here, you can identify the servers to which the user is connected and determine which should and should not be optimized. Figure 2-7 shows an example of the output of this command.

```

C:\Documents and Settings\Administrator.PEAP>netstat -n

Active Connections

Proto Local Address          Foreign Address        State
TCP   10.10.13.100:3389      10.21.113.227:53323   ESTABLISHED
TCP   10.10.13.100:4493     10.10.10.100:445     TIME_WAIT
TCP   10.10.13.100:4546     10.10.10.100:445     ESTABLISHED
TCP   10.10.13.100:4550     10.10.10.100:21      TIME_WAIT
TCP   10.10.13.100:4553     10.10.10.100:21      TIME_WAIT
TCP   10.10.13.100:4555     10.10.10.100:21      TIME_WAIT
TCP   10.10.13.100:4557     10.10.10.100:21      TIME_WAIT
TCP   10.10.13.100:4562     10.10.10.100:21      TIME_WAIT
TCP   10.10.13.100:4565     10.10.10.100:21      TIME_WAIT
TCP   10.10.13.100:4567     10.10.10.100:21      TIME_WAIT
TCP   10.10.13.100:4569     10.10.10.100:21      TIME_WAIT
TCP   10.10.13.100:4574     10.10.10.100:21      TIME_WAIT
TCP   10.10.13.100:4577     10.10.10.100:21      TIME_WAIT
TCP   10.10.13.100:4579     10.10.10.100:21      TIME_WAIT
TCP   10.10.13.100:4581     10.10.10.100:21      TIME_WAIT
TCP   10.10.13.100:4585     10.10.10.100:21      ESTABLISHED
TCP   127.0.0.1:5152        127.0.0.1:4219      CLOSE_WAIT

```

Figure 2-7 Determining the Number of TCP Connections In Use Using *netstat*

Another tool provided by Microsoft that (along with many other things) provides visibility into the number of TCP connections in use on a particular computer is Performance Monitor. Performance Monitor can be accessed by clicking **Start > Run** and typing **perfmon**, followed by clicking **Ok**. From within the Performance Monitor window, click the **+** sign, select the TCP performance object, and then add the **Connections Established** counter. Doing so shows you the number of connections established over time, and this data can even be exported for offline use. Figure 2-8 illustrates an example output from Performance Monitor showing the number of established TCP connections.

Linux, UNIX, and Macintosh provide similar tools to understand the number of connections that are open on a given computer. The **netstat** command is available on virtually any Linux distribution and is available in most UNIX platforms and versions of Apple's Macintosh OS/X operating system.

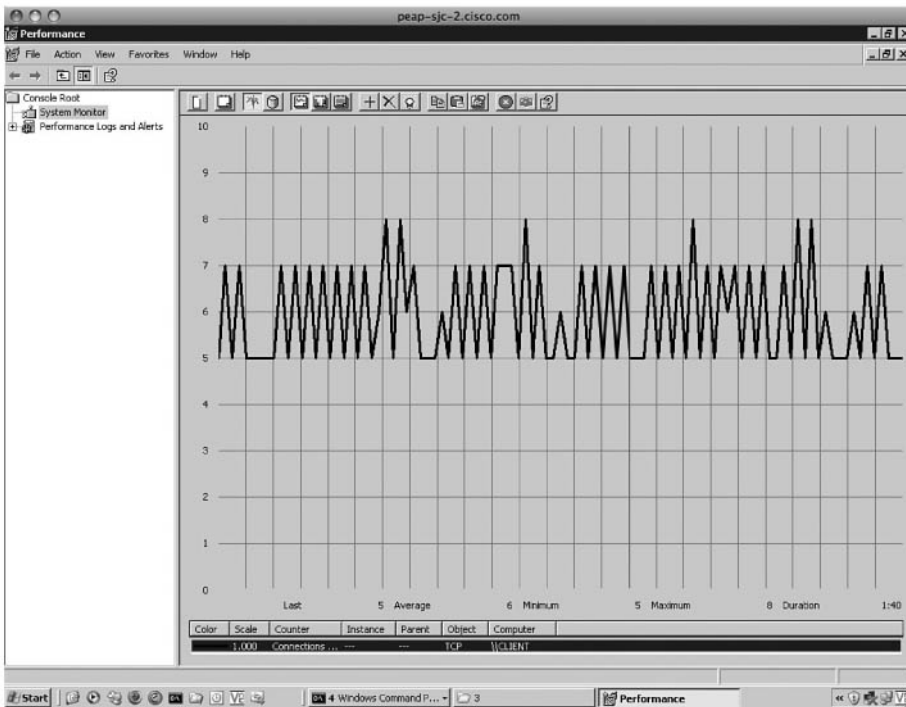


Figure 2-8 Determining the Number of TCP Connections in Use Using Performance Monitor

For the data center, the sum of all remote office TCP connections should be considered one of the key benchmarks by which the data center sizing should be done. Note that the largest Cisco WAAS device supports up to 50,000 optimized TCP connections—which is approximately 5,000 users (assuming ten TCP connections per user). For organizations that need to support a larger number of users or want to deploy the data center devices in a high-availability manner, multiple devices can be used. The type of network interception used (discussed in Chapter 4) determines the aggregate number of optimized TCP connections that can be supported by a group of Cisco WAAS devices deployed at a common place within the data center. Recommended practice dictates that sites that require high availability be designed with $N+1$ availability in consideration relative to the number of maximum optimized TCP connections—that is, if 100,000 optimized TCP connections must be supported, the location should have a minimum of two WAE-7371 devices to support the workload, a third WAE-7371 device to handle failure of one of the devices, and use an interception mechanism such as WCCP or ACE that supports load-balancing of workload across the entire set of three devices. Other considerations apply, as discussed in Chapter 4.

WAN Bandwidth and LAN Throughput

WAAS devices are not restricted in software or hardware in terms of the amount of WAN bandwidth or LAN throughput supported. However, recommendations are in place to specify which WAAS device should be considered for a specific WAN environment. WAN bandwidth is defined as the amount of WAN capacity that the WAAS device can fully use when employing the full suite of optimization capabilities (this includes DRE, PLZ, TFO, and the other application acceleration capabilities). LAN throughput is defined as the maximum amount of application layer throughput (throughput as perceived by the users and servers) that can be achieved with the particular WAAS hardware model and an equivalent or more-powerful peer deployed at the opposite end of the network.

For some deployment scenarios, it is desired to use the Cisco WAAS devices only for TCP optimization. Cisco WAAS TFO provides a powerful suite of optimizations to better allow communicating nodes to “fill the pipe” (that is, fully leverage the available WAN bandwidth capacity) when the application protocol is not restricting throughput due to application-induced latency. Each Cisco WAAS device has a TFO-only throughput capacity that can be considered when WAAS devices are deployed strictly for TCP optimization only. This is recommended only for situations where compression, redundancy elimination, and application acceleration are not required, and the application throughput has been validated to be hindered only by the performance of the TCP implementation in use. This is common in some data center to data center applications—such as data replication or data protection—where the traffic that is sent is previously compressed, redundancy eliminated, or encrypted. TFO attempts to fully utilize the available bandwidth capacity, but might be hindered by congestion in the network (not enough available bandwidth) or performance impedance caused by application protocol chatter.

Table 2-4 shows the WAN bandwidth supported by each WAAS device model and the maximum LAN-side throughput and TFO-only throughput capacity. Note that other factors can influence these values and throughput levels can be achieved only when the link capacity available supports such a throughput level. For instance, a LAN throughput maximum of 150 Mbps is not possible on a Fast Ethernet connection; rather, a Gigabit Ethernet connection is required. Similarly for throughput speeds more than 1 Gbps, multiple 1-Gbps interfaces must be used.

The amount of bandwidth required per site is the sum of available WAN capacity that can be used at that site and not the sum of all WAN bandwidth for every connected peer. For instance, if a branch office has four bundled T1 links (totaling 6 Mbps of aggregate WAN throughput) but only two are used at any given time (high availability configuration), a device that supports 3 Mbps or more is sufficient to support the location.

Table 2-4 WAN Bandwidth and LAN Throughput Capacity per WAAS Device

WAAS Device Model	WAN Supported	LAN Throughput Maximum	TFO-Only Throughput Maximum
NME-WAE-302	4 Mbps	90 Mbps	100 Mbps
NME-WAE-502	4 Mbps	150 Mbps	150 Mbps
NME-WAE-522	8 Mbps	2000 Mbps	250 Mbps
WAVE-274	2 Mbps	90 Mbps	150 Mbps
WAVE-474	4 Mbps	90 Mbps	250 Mbps
WAE-512-1GB	8 Mbps	100 Mbps	350 Mbps
WAE-512-2GB	20 Mbps	150 Mbps	400 Mbps
WAVE-574-3GB	8 Mbps	100 Mbps	350 Mbps
WAVE-574-6GB	20 Mbps	150 Mbps	400 Mbps
WAE-612-2GB	45 Mbps	250 Mbps	450 Mbps
WAE-612-4GB	90 Mbps	350 Mbps	500 Mbps
WAVE-674-4GB	45 Mbps	250 Mbps	450 Mbps
WAVE-674-8GB (with or without VB)	90 Mbps	350 Mbps	500 Mbps
WAE-7326	155 Mbps	450 Mbps	600 Mbps
WAE-7341	310 Mbps	800 Mbps	800 Mbps
WAE-7371	1 Gbps	1.5 Gbps	1.8 Gbps

Similarly, if a data center has four DS-3 links (totaling 180 Mbps of aggregate WAN throughput) but uses only three at a time ($N+1$ configuration), a device that supports 135 Mbps of WAN bandwidth or more is sufficient to support that location. The WAN throughput figures mentioned in the preceding table are (as discussed previously) not limited in hardware or software. In some cases, the WAN throughput that a device achieves might be higher than the values specified here. Those interested in using a smaller device to support a larger WAN link (for instance, qualifying a WAVE-274 for an 8-Mbps ADSL connection) are encouraged to test the system under those conditions and validate the performance prior to making a decision to use that specific platform.

Number of Peers and Fan-Out

Each Cisco WAAS device has a static system limit in terms of the number of concurrent peers it can actively communicate with at any one given time. When designing for a particular location where the number of peers exceeds the maximum capacity of an individual device, multiple devices can be deployed, assuming an interception mechanism that uses load balancing is employed (such as WCCPv2 or ACE; these are discussed in

Chapter 4). In cases where load balancing is used, TCP connections are distributed according to the interception configuration, thereby allowing for near-linear scalability increases in connection count, peer count, and WAN bandwidth, as devices are added to the pool. Load-balancing interception techniques are recommended when multiple devices are used in a location, and in general, an $N+1$ design is recommended.

Peer relationships are established between Cisco WAAS devices during the automatic discovery process on the first connection optimized between the two devices. These peer relationships time out after ten minutes of inactivity (that is, no active connections are established and optimized between two peers for ten minutes). Each WAAS device supports a finite number of active peers, and when the peer relationship is timed out, that frees up peering capacity that can be reused by another peer. Data stored in the DRE compression history remains intact even if a peer becomes disconnected due to inactivity, unless the DRE compression history becomes full. In cases where the DRE compression history becomes full, an eviction process is initiated to remove the oldest set of data in the DRE compression history to make room for new data.

Table 2-5 shows the maximum number of concurrent peers supported per WAAS platform. If peers are connected beyond the allocated limit, the WAE permits the connections to be

Table 2-5 *Maximum Supported Peers per WAAS Device*

Network Module	Concurrent Peers	Appliance	Recommended Concurrent Peers
302	5	WAVE-274	35
502	15	WAVE-474	35
522	40	512-1GB	35
		512-2GB	70
		WAVE-574-3GB	35
		WAVE-574-6GB	70
		612-2GB	210
		612-4GB	350
		WAVE-674-4GB	100
		WAVE-674-8GB (no VB)	200
		WAVE-674-8GB (with VB)	200
		7326	600
7341	1400		
7371	2800		

established and gracefully degrades performance as needed. Connections associated with peers in excess of the maximum fan-out ratio are able to use the existing compression history but are not able to add new chunks of data to it. The end result is lower effective compression ratios for the connections using peers that are in excess of the specified fan-out ratio.

The number of peers supported by a device is typically the last factor that should be considered when sizing a solution for a particular location. The primary reason being that the WAN capacity or number of connections supported at the maximum concurrent peers specification is generally an order of magnitude higher than what the device can support. For instance, although a WAE-7371 can support up to 2800 peers, even if those peers were the NME-302 (each supporting 250 optimized TCP connections), it is not able to handle the 700,000 possible optimized TCP connections that all 2,800 NME-302s were attempting to optimize with it. It is best to size a location first based on WAN bandwidth capacity and TCP connections, and in most cases, only a simple validation that the number of peers supported is actually required.

Number of Devices Managed

Each Cisco WAAS deployment must have at least one Cisco WAAS device deployed as a Central Manager. The Central Manager is responsible for system-wide policy definition, synchronization of configuration, device monitoring, alarming, and reporting. The Central Manager can be deployed only on appliances and can be deployed in an active/standby fashion. When a certain WAAS device is configured as a Central Manager, it is able to, based on the hardware platform selected for the Central Manager, manage a maximum number of WAAS devices within the topology. Only WAAS appliances can be configured as Central Manager devices, and in high-availability configurations, each Central Manager WAE should be of the same hardware configuration. Although hardware disparity between Central Manager WAEs works, it is not a recommended practice given the difference in the number of devices that can be managed among the WAE hardware models. It should be noted that standby Central Managers (such a configuration is examined in Chapter 7) receive information in a synchronized manner identical to how accelerator WAAS devices do. Table 2-6 shows the maximum number of managed nodes that can be supported by each WAAS appliance when configured as a Central Manager.

Use of multiple WAAS devices configured as Central Manager devices do not increase the overall scalability in terms of the number of devices that can be managed. To manage a number of devices greater than the capacities mentioned in the preceding table, multiple autonomous Central Managers are needed. For instance, in an environment with 3000 devices, two separate instances of Central Manager are required, and each instance can be comprised of a single device or multiple devices deployed in a high availability primary/standby configuration.

Table 2-6 *Central Manager Scalability*

Appliance	Managed Nodes
WAVE-274	125
WAVE-474	250
WAE-512-1GB	500
WAE-512-2GB	750
WAVE-574-3GB	500
WAVE-574-6GB	1000
WAE-612-2GB	750
WAE-612-4GB	1500
WAE-674-4GB	1500
WAE-674-8GB	2000

Replication Acceleration

The WAE-7341 and WAE-7371 devices support a deployment mode called *Replication Accelerator*, which requires Cisco WAAS version 4.0.19, or a version newer than that from the 4.0 train. This mode of acceleration is used for data center to data center deployments where replication and backup acceleration is required, and when configured, adjusts the behavior of the WAAS device to allocate larger blocks of memory to a smaller number of connections, and minimizes the processing latency of DRE by using only memory for deduplication. Although only memory is used for DRE, the DRE data is persistent in that it is written to disk, but the disk is used only to reload the previous compression history. This enables WAAS to provide high levels of throughput necessary to accelerate replication and backup traffic between data centers.

The network typically found in these cases is high-bandwidth and relatively low latency (above 10–20 ms), where a significant amount of data needs to be moved from one location to another location in a short period of time. The performance and scalability metrics of replication accelerator mode are different than the performance and scalability metrics that would normally be considered for these devices when not deployed in replication accelerator mode and are documented in Table 2-7.

Table 2-7 *Replication Accelerator Performance and Scalability Metrics*

Appliance	WAN Bandwidth	LAN Throughput	Optimized TCP Connections	Concurrent Peers	DRE Capacity
WAE-7341	310 Mbps	800 Mbps	2500	4	12 GB
WAE-7371	1 Gbps	1.5 Gbps	5000	9	24 GB

Although all WAAS devices in a given network can be managed by a common Central Manager, WAAS devices configured in replication accelerator mode can only peer with other WAAS devices that are configured as replicator accelerator devices. Should intermediary application accelerator devices exist in the network path between two replication accelerator devices (this is generally rare, as replication accelerator devices are deployed between backend networks as opposed to the enterprise WAN), the application accelerator devices are not able to peer with replication accelerator devices.

Replication accelerator devices are commonly deployed on backend data center to data center networks and not the enterprise WAN due to the high bandwidth requirements. WAAS devices configured as replication accelerators are commonly found deployed as follows:

- **Directly attached to one or more storage array IP/Ethernet interfaces:** Such a deployment model dedicates the devices to optimize replication for that particular array and that particular interface.
- **Directly attached to one or more storage fabric switch or director IP/Ethernet interfaces:** Including the Cisco MDS 9000 family, such a deployment model enables the devices to optimize replication or backup traffic traversing fabrics in distant sites over IP.
- **Directly behind the data center interconnect device:** Such a deployment model enables optimization of any traffic between data centers. In this deployment model, replication accelerator should be carefully considered against the standard application accelerator mode which may be more applicable in cases where a large body of non-replication and nonbackup traffic exists.

Virtual Blades

The Cisco WAVE appliance family and the WAE-674 provide branch office virtualization capabilities that enable consolidation of remote branch office servers onto the WAAS device as a shared platform. Sizing for VBs should be done in conjunction with sizing for WAN optimization and application acceleration because the available disk capacity to support VBs and the number of VBs supported varies per platform based on the hardware configuration as shown in Table 2-8.

To accurately size a virtualization solution for a branch office, it is necessary to understand the minimum and recommended memory requirements to support the operating system and applications you plan to install on top of that operating system. Many vendors support installation of their server operating system onto systems with only 512 MB of memory, which increases the maximum number of VBs that can be installed on a WAAS device; however, many have requirements for larger amounts of memory.

Additionally, consider the disk capacity requirements necessary for each VB, and reconcile that amount with the total VB storage capacity of the platform selected for that given location. Even the smallest virtualization-capable WAAS device (the WAVE-274) supports 35 GB of disk capacity for VBs—meaning that with two VBs, configured, you have

Table 2-8 *VB Capacity*

Appliance	VB Disk Capacity	VB Memory Capacity	Maximum Number of VBs (512 MB RAM each)
WAVE-274	35 GB	1 GB	2
WAVE-474	35 GB	1 GB	2
WAVE-574-3GB	60 GB	1 GB	2
WAVE-574-6GB	180 GB	3 GB	6
WAE-674-4GB	120 GB	1 GB	2
WAE-674-8GB	200 GB	3 GB	6

approximately 17.5 GB of disk space for each. Storage capacity allocation is flexible in that you can allocate as much space as is available from the pool to any particular VB. However, you should ensure that you size the system for the location with enough capacity to support the current application and operating system requirements as well as future requirements. More information on configuration and deployment of VBs can be found in Chapter 10, “Branch Office Virtualization.”

Summary

The Cisco Wide-Area Application Engine family includes three network modules for the Integrated Services Router and six appliance models spanning two desktop models and four rack-mount appliance models. This breadth of portfolio provides customers with the flexibility necessary to allocate the right platform for each network location where WAN optimization, application acceleration, and virtualization capabilities are needed. Four licenses are available for Cisco WAAS, including the Transport license (WAN optimization capabilities only), Enterprise license (all application accelerators except video, and certain other features), Video (Windows Media over RTSP stream splitting), and Virtual-Blades (branch office virtualization platform). Sizing of a Cisco WAAS solution requires consideration of a number of factors, including network conditions (WAN bandwidth and LAN throughput), number of users and concurrent optimized TCP connections, disk capacity and compression history, memory, concurrently connected peers, and virtualization requirements. By following the recommended guidelines for performance and scalability, a robust Cisco WAAS design can be realized, thereby allowing administrators to deploy the solution confidently to improve application performance over the WAN while enabling centralization and consolidation of costly infrastructure.

Index

A

- AAA Accounting, 103
- ABE (Microsoft Access-Based Enumeration), 92
- Accelerate parameter (policy maps), 359
- accelerated replication, 74–75
- Acceleration tab (monitoring charts), 292
- acceleration, application, 401–403
 - charts, 460–461, 566
 - CIFS acceleration, 403–406
 - enabling, 415–416, 419–423
 - HTTP acceleration, 411–412
 - MAPI acceleration, 409–411
 - monitoring charts, 291–295
 - monitoring with CM GUI, 460–462
 - monitoring with device CLI, 453–459
 - monitoring with XML-API, 463
 - CIFSStats service*, 463–465
 - HrtpStats service*, 467–468
 - MapiStats service*, 468–469
 - NfsStats service*, 470
 - SSLStats service*, 466–467
 - VideoStats service*, 467
 - NFS acceleration, 408–409
 - SSL acceleration, 412–414, 432–447
 - video acceleration, 414–415, 423–425
 - Windows print acceleration, 407–408
- Access-Based Enumeration (ABE), 92
- accessible data, XML-API integration, 310
 - alarm status service, 311–312
 - device configuration service, 310–311
 - device status service, 312
 - traffic acceleration service, 312–313
- ACE (Application Control Engine), 144–145, 210, 227
 - bridged mode deployment, 227–230
 - routed mode deployment, 230–232
 - scalability, 239–240
- ACK (acknowledge) packets, 325–326
- ACNS (Application and Content Networking Services), 106
- Action parameter (policy maps), 359
- activation
 - devices, 270–271
 - licenses, 478
- AD (Active Directory), 106, 280
- Adaptive Security Appliance (ASA), 335
- adaptive TCP buffering, configuring WAN optimization, 339–345
- Add/Edit Interface window, 487
- addresses
 - IP
 - multiple next-hop addresses*, 198
 - verifying next-hop addresses*, 197
 - tunnel interface configuration, 147
 - VIP, 146
- admin drawer, CM My WAN context homepage, 264
- administrative groups, Windows parameters, 282

- Advanced Technology Attachment (ATA), 486
- AES-256 encryption, 50
- aggregation
 - link, WAE, 111–115
 - payload, 410
- Alarm Book, 102, 290
- alarm status service (XML-API), 311–312
- alarms, 290
 - admin drawer, 264
 - management, 290–291
 - monitoring virtual blades, 505–506
 - remediation actions, 291
 - services, Cisco Linux platform, 52
- AllDevicesGroup device group, 272
- analysis, WAAS requirements for deployment, 78–79
 - application characteristics, 90–91
 - Application Optimizer (AO) requirements, 91–98
 - availability requirements, 99–100
 - management requirements, 100–103
 - network infrastructure, 82–90
 - platform requirements, 98–99
 - scalability requirements, 99
 - security requirements, 103–105
 - site information, 80–82
 - virtualization requirements, 105–106
- AO (Application Optimizer), 331, 401
 - advanced features, 92
 - CIFS AO, 91–92
 - enabling, 416, 419–420, 423
 - file services utilization, 93
 - HTTP AO, 95–96
 - MAPI AO, 94–95
 - NFS AO, 96
 - Replication Accelerator, 98
 - SSL AO, 97
 - Video AO, 96–97
- APM (Application Performance Management), 56, 393
- appliances, 58–59
 - hardware architecture, 49
 - WAE model 512, 60
 - WAE model 612, 60
 - WAE model 674, 61
 - WAE model 7341, 61
 - WAE model 7371, 61
 - WAVE model 274, 59
- WAVE model 474, 59
- WAVE model 574, 60
- application acceleration, 401–403
 - charts, 460–461
 - CIFS acceleration, 403–406
 - enabling, 415–416, 419–420, 423
 - HTTP acceleration, 411–412
 - MAPI acceleration, 409–411
 - monitoring using CM GUI, 460, 462
 - monitoring using device CLI, 453–459
 - monitoring using XML-API, 463
 - CIFSStats service*, 463–465
 - HttpStats service*, 467–468
 - MapiStats service*, 468–469
 - NfsStats service*, 470
 - SSLStats service*, 466–467
 - VideoStats service*, 467
 - NFS acceleration, 408–409
 - SSL acceleration, 412–414, 432–447
 - video acceleration, 414–415, 423–425
 - Windows print acceleration, 407–408
- Application and Content Networking Services (ACNS), 106
- Application Classifier parameter (policy maps), 359
- Application Control Engine. *See* ACE
- application groups (ATP), 348–352
- Application Optimizer. *See* AO
- Application parameter (policy maps), 359
- Application Performance Management (APM), 56, 393
- Application Requirements Checklist, 91
- Application Response Time monitoring. *See* ART
- application traffic policy. *See* ATP
- application-accelerator device mode, 257
- applications
 - acceleration. *See* acceleration
 - performance, 3
 - CM optimization results*, 565–566
 - hot transfer measurements of FTP*, 565
 - measuring*, 564
 - testing tools*, 566
 - Windows File Transfer*, 564–565
 - requirements for deployment, 90–91
 - Application Optimizer (AO) requirements*, 91–98
 - Application Requirements Checklist*, 91
 - availability requirements*, 99–100

- management requirements, 100–103*
- platform requirements, 98–99*
- scalability requirements, 99*
- security requirements, 103–105*
- virtualization requirements, 105–106*

applied policies, 574

architecture

- hardware, 49, 55
 - appliances, 58–61*
 - licenses, 61–62*
 - router-integrated network modules, 56–58*
- per-peer context, 323
- software, 50–55

ART (Application Response Time) monitoring, 394–399

ASA (Adaptive Security Appliance), 335

ASCH Password Authentication setting (ITACACS+ configuration), 286

asymmetric routing, dual data centers, 224–226

asymmetric traffic routing, 207–208

ATA (Advanced Technology Attachment), 486

ATP (Application Traffic Policy), 94, 347–348

- application groups, 348, 351–352
- Central Manager configurations for custom policies, 544–545

- EndPoint Mapper classification, 366–370

- policy maps, 358–365

- policy negotiation, 365–366

- traffic classifiers, 352, 355–357

- WAAS architecture, 348

ATP (Application Traffic Policy engine) engine, 50–52

- Cisco Linux platform, 53–54

attributes, WCCP service groups, 121–123

audit trails, 296–297

authentication

- Kerberos, 282
- nonlocal users, 278–279
- NTLM, 282
- provisioned device management, 280–289
- settings configuration, 279

authentication configuration commands, 286

authentication login commands, 286

Authentication Status window, 283

automatic discovery, 370–372, 574

automatic discovery mechanism, 324–327

AutoStart parameter (VB configuration), 485

availability, requirements for deployment, 99–100

Availability Checklist, 100

B

backup

- CM database, 305–307

- virtual blades, 501–502

bandwidth

- scalability (TFO optimization), 322

- WAE, 111–115

- WANs, 70–71

bandwidth command, 109

bandwidth delay product (BDP), 341

baseline performance tests, 549

- FTP measurements, 551

- Windows file services, 549–551

Basic setting (policy map Type parameter), 358

BDP (bandwidth delay product), 341

best practices, network integration, 150–151

BIC-TCP (Binary Increase Congestion TCP), 340

Binary Increase Congestion TCP (BIC-TCP), 340

blacklist operation (TFO)

- configuration, 336

- configuration of WAN optimization, 333–337

blade number parameter (VB configuration), 485

Blue Screen of Death, 507

Boot from parameter (VB configuration), 485

boot sequence (devices), 256, 497–500

branch office virtualization, 473

- business benefits, 474

- overview, 473–475

- virtual blades, 475–476

- accessing console, 495–496*

- changing boot sequence, 497–500*

- creating virtual blades, 478–493*

- hardware emulation, 476–477*

- management, 476, 500–502*

- monitoring, 503–506*

- platforms and capacity, 477*

- starting virtual blades, 493–494*

- stopping, 496–497*

- troubleshooting, 506–509*

branch offices

- nonredundant, 154–158, 163–167, 175–181

- redundant, 158–161, 181–186, 189–191, 194, 196

- serial inline clustering, 162–163

bridged mode (ACE), 227–230

BSOD (Blue Screen of Death), 507

buffering

- TCP, WAN optimization, 339–345
- TFO configuration, 342, 344

built-in setup wizard, 250–252, 256–260**business benefits, 474****bypass manager, 52****C****caching**

metadata

- CIFS acceleration, 404*
- Windows print acceleration, 407*

safe data, 404

video-on-demand, 415

CAD/CAM (Computer Aided Design/Computer Aided Manufacturing), 408**capacity**

disks, 65

number of devices managed, 73

peers and fan-out, 71–73

Replication Accelerator, 74–75

TCP connections, 65–69

VBs (virtual blades), 477

virtual blades, 75, 477

WAN bandwidth and LAN throughput, 70–71

case studies

Data Center

- interception method, 533–534*
- network topology, 533*
- requirements, 533*
- WAE configuration, 534, 536*
- WAE placement, 533–534*
- WAN router configuration, 537, 540, 543*

IData CenterA, 532

remote site profile A, 512

- interception method, 513*
- LAN switch configuration, 517–519*
- network topology, 513*
- requirements, 513*
- WAE configuration, 513–515*
- WAE placement, 513*
- WAN router configuration, 516–517*

remote site profile B, 519

- interception method, 520*
- network topology, 520*
- requirements, 519*

*WAE configuration, 520–522**WAE placement, 520**WAN router configuration, 522–524, 528–532*

remote site profile C, 524

- interception method, 525–526*
- network topology, 525*
- requirements, 524*
- WAE configuration, 526–528*
- WAE placement, 525–526*

requirements, 511

WAN topology, 511–512

Catalyst 6500 Series switches, 203**CBQoS (class-based QoS), 381****CD-Image parameter (VB configuration), 485****CDN (Content Delivery Network), 96****Central Manager. See CM****Centralized Management System. See CMS****charts**

- acceleration, 460–461
- monitoring, 291–295

checklists, requirements for deployment

- Application Requirements Checklist, 91
- Availability Checklist, 100
- File Services Requirements Checklist, 93
- HTTP Requirements Checklist, 95–96
- Management Requirements Checklist, 103
- MAPI Requirements Checklist, 95
- Network Infrastructure Checklist, 89–90
- NFS Requirements Checklist, 96
- Platform Requirements Checklist, 98–99
- Security Requirements Checklist, 105
- Site Information Checklist, 82
- SSL Requirements Checklist, 97
- Video Requirements Checklist, 96–97
- Virtualization Requirements Checklist, 106

CIFS

- acceleration, 403–406
- cache capacity, 65
- cold transfers, 564
- ports, 262

CIFS AO (CIFS Application Optimizer), 91–92**CIFS setting (policy map Accelerate parameter), 359****CIFSStats service, 463–465****Cisco IOS Firewall (IOS FW), 199–201****Cisco Linux platform, 50**

- ATP engine, 53–54
- CMS (Central Management Subsystem), 51

- data encryption, 50–51
- interface managers, 51
- monitoring facilities and alarms, 52
- network interception and bypass manager, 52
- virtual blades, 55
- Cisco Port Aggregation Protocol.** *See* PAgP
- Cisco Technical Assistance Center (TAC),** 592
- Cisco WAAS Mobile,** 3
- class-based QoS (CBQoS),** 381
- clear cache command,** 565
- CLI (command-line interface),** 50, 260–261
 - acceleration monitoring, 453–459
 - commands, 595–597
 - configuration, 239–240, 287
 - application groups,* 349
 - SSL acceleration services,* 438–447
 - CPU utilization monitoring, 388
 - disk monitoring, 389
 - DRE monitoring, 392
 - stopping a virtual blade, 497
 - VB resource allocation, 476
- client distribution, WCCP scalability,** 234–239
- client-side WAE, configuring with Setup Wizard,** 559–563
- clustering, serial inline clustering,** 162–163
- CM (Central Manager)**
 - backing up virtual blade files, 501
 - configuration, 258, 289, 544–545, 552–555
 - alarms and monitors,* 290–295
 - backup and restoration of database,* 305–307
 - logging mechanisms,* 296–301
 - programmatic interfaces,* 308–316
 - reports,* 295–296
 - software upgrades/downgrades,* 302–305
 - SSL acceleration services,* 433–434, 437
 - deployment requirements, 100
 - domain definitions, 104
 - enabling services manually, 259
 - encryption keys, 50
 - GUI. *See* CM GUI, 250
 - GUI login page, 262
 - key management, 261
 - My WAN context homepage, 262–265
 - optimization, 565–566
 - overview, 261–266
 - registration, verification, 563
 - role definitions, 104
 - scalability, 73, 262
 - security, 261
 - simplicity, 262
 - specification of hostname, 259
 - stopping a virtual blade, 497
 - system timers, 267
 - virtualization management roles, 500
- CM GUI**
 - configuring application groups, 351
 - Connection Statistics page, 378
 - CPU utilization monitoring, 389
 - FlowAgent configuration, 399
- CMS (Central Management Subsystem),** 50, 250, 266–268
 - Cisco Linux platform, 51
 - database backup, 306
 - database restoration, 307
 - downgrading the database, 303
 - registration and service status, 268
- cms database restore command,** 307
- cms enable command,** 260, 266
- cms recover identity command,** 269
- cold transfers,** 564
- collection of requirements**
 - deployment, 78–79
 - application characteristics,* 90–91
 - Application Optimizer (AO) requirements,* 91–98
 - availability requirements,* 99–100
 - management requirements,* 100–103
 - network infrastructure,* 82–90
 - platform requirements,* 98–99
 - scalability requirements,* 99
 - security requirements,* 103–105
 - site information,* 80–82
 - virtualization requirements,* 105–106
- command-line interface.** *See* CLI, 50, 260
- commands**
 - authentication configuration, 286
 - authentication login, 286
 - bandwidth, 109
 - clear cache, 565
 - CLI, 595, 597
 - cms database restore, 307
 - cms enable, 260, 266
 - cms recover identity, 269
 - copy cdrom wow-recovery, 483

- copy disk ftp, 300
 - copy running-config startup-config, 253
 - copy sysreport {diskl ftp | tftp}, 592
 - dir, 300, 480
 - inline vlan all, 161
 - inline vlan number, 579
 - inspect waas, 243, 584
 - ip inspect WAAS enable, 584
 - ip name-server, 283
 - ip wccp redirect exclude in, 134, 180
 - ip wccpservice groupaccelerated, 581
 - mac-sticky, 228
 - map, 362
 - map basic disable (#), 362
 - map basic insert, 362
 - map basic list, 362
 - map basic move from (#) to (#), 362
 - mkdir, 300
 - net use, 550
 - next-hop verify-availability route map, 197
 - no normalization, 230
 - no vnc, 495
 - no wccp ver 2, 128
 - policy-engine application, 362
 - service-module ip addressaddrmask, 170
 - service-module ip default-gatewayaddr, 171
 - set ip next-hop, 139
 - setip next-hop, 198
 - setup, 253
 - show, 343
 - show accelerator, 419, 587
 - show auto-discovery blacklist, 372
 - show auto-discovery list, 372
 - show clock, 283
 - show cms info, 267, 563
 - show conn longx, 240
 - show disk details, 480
 - show disks details, 591
 - show hardware, 577
 - show interface, 110–111, 149, 256, 572
 - show interface inlinegroup slot/number, 578
 - show interface PortChannel, 113–115
 - show interface Standby, 117
 - show ip wccp, 580
 - show ip wccpservice groupdetail, 579
 - show license, 478
 - show running-config, 336, 577
 - show statistics accelerator CIFS details, 588
 - show statistics application, 383
 - show statistics auto-discovery, 370
 - show statistics auto-discovery blacklist, 371
 - show statistics connection, 373–375, 454, 588
 - show statistics connection conn-idid, 585
 - show statistics connection optimized, 588
 - show statistics connection pass-through, 574
 - show statistics dre, 392
 - show statistics pass-through, 575
 - show tcam interface, 581
 - show tech-support command, 593
 - show virtual-blade, 481, 503
 - show virtual-bladenumblockio, 504
 - show virtual-bladenumberinterface, 504
 - show wccp gre, 583
 - show wccp routers, 582
 - show wccp services, 582
 - show wccp status, 582
 - showstatistics acceleratoracceleratordetail, 457
 - snmp-server, 300
 - ssh-key-generate, 260
 - standby, 109
 - test self-diagnostictest, 590
 - type, 300
 - type-tail, 300
 - type-tail /local1/errorlog/virtual-blade, 506
 - virtual-blade n start, 494
 - virtual-bladenumbercd eject, 498
 - virtual-bladenumberkill-save-state, 506
 - virtual-bladenumbersave, 502
 - wccp tcp-promiscuous service group, 170
 - write memory, 253
- Comments property (domain configuration page), 276**
- community strings (SNMP), 101–102**
- compression**
- e-mail, 410
 - history, 64
 - PLZ (Persistent LZ Compression), 324
 - scalability, 323
- Computer Aided Design/Computer Aided Manufacturing. See CAD/CAM**
- configure drawer, CM My WAN context homepage, 264**
- connection statistics, 373–374, 378–380**
- connection reuse, 407**
- Connection Statistics page (CM GUI), 378**

connections

- EMIC, 145–149
- optimized, 320–321
- original, 320
- setup time, 395

connectivity

- NME-WAE interfaces, 108
- WAE interfaces, 107–111
 - link aggregation*, 111–115
 - standby interface feature*, 115–119

console (VBs), accessing, 495–496

Content Delivery Network. *See* CDN

content switching, 143–145

contexts (My WAN), changing, 264–265

copy cdrom wow-recovery command, 483

copy disk ftp command, 300

copy running-config startup-config command, 253

copy sysreport [diskl ftp | tftp] command, 592

core locations, repositioning, 447–449

counters

- Transparent GRE packets received, 583
- Transparent non-GRE non-WCCP packets received, 583
- Transparent non-GRE packets received, 583

CPUs

- dedicated, 476
- emulation parameter (VB configuration), 487
- utilization monitoring, 388–389

critical alarms, VB faulty shutdown, 505

custom policies, Central Manager configuration, 544–545

D**data, XML-API services, 310**

- accessing with soapUI tool, 313–316
- alarm status service, 311–312
- device configuration service, 310–311
- device status service, 312
- traffic acceleration service, 312–313

Data Center case study, 532

- interception method, 533–534
- network topology, 533
- requirements, 533
- WAE configuration, 534, 536
- WAE placement, 533–534
- WAN router configuration, 537, 540, 543

data center network integration, 203

- data center topology, 203
 - asymmetric traffic flows*, 207–208
 - dual data centers*, 205
 - multi-data centers*, 205
 - server farm distribution*, 209–212
 - single WCCP service groups*, 209
 - WAAS placement*, 205
 - WAN edge*, 205
- deployment solutions, 212
 - server load balancing*, 227–228, 232
 - WCCP*, 212–213, 216–226
- firewalls, 240–241, 243–246
 - FWSM connection display output*, 240–241
 - PIX/ASA configuration*, 243–246
 - server farm aggregation with FWSM*, 241–243
- scaling transparent interception, 233
 - ACE*, 239–240
 - WCCP*, 233–239

data center topology, 86**data encryption, 50–51****data feed poll rate, 266****data reduction, 415****data transfer time, 395****database backup and restoration, 305–307****dedicated CPU(s), 476****default policies, restoring, 364****defining**

- application groups, 349
- traffic classifiers, 353

degradation, troubleshooting, 570

- application acceleration, 587–589
- firewall integration, 584
- half-duplex conditions, 572–573
- interception, 577–583
- low-compression ratios, 584–587
- pass-through traffic, 573–576

deployment

- application characteristics, 90–91
- Application Optimizer (AO) requirements, 91
 - advanced features*, 92
 - CIFS AO*, 91–92
 - file services utilization*, 93
 - HTTP AO*, 95–96
 - MAPI AO*, 94–95
 - NFS AO*, 96
 - Replication Accelerator*, 98

- SSL AO, 97
- Video AO, 96–97
- availability requirements, 99–100
- data center network integration, 203, 212
 - data center topology*, 203–212
 - firewalls*, 240–246
 - scaling transparent interception*, 233–240
 - server load balancing*, 227–228, 232
 - WCCP*, 212–213, 216–226
- in-path, 153–154
 - nonredundant branch offices*, 154–158
 - redundant branch offices*, 158–161
 - serial inline clustering*, 162–163
- management requirements, 100
 - CM and XML-API*, 100
 - Management Requirements Checklist*, 103
 - SNMP community strings*, 101–102
 - SNMP traps/informs*, 101
 - syslog servers*, 102
- network infrastructure, 82
 - data center topology*, 86
 - Network Infrastructure Checklist*, 89–90
 - remote office topology*, 85–86
 - traffic flows*, 87–89
 - WAN topology*, 82–84
- off-path, 163
 - IOS FW*, 199–201
 - nonredundant branch offices*, 163–181
 - policy-based routing interception*, 196–199
 - redundant branch offices*, 181–186, 189–191, 194, 196
- planning overview, 77–78
- platform requirements, 98–99
- Replication Accelerator, 74
- requirements collection and analysis, 78–79
- scalability requirements, 99
- security requirements, 103–105
- site information, 80–82
- troubleshooting, 570–571
 - application acceleration*, 587–589
 - firewall integration*, 584
 - half-duplex conditions*, 572–573
 - interception*, 577–583
 - low-compression ratios*, 584–587
 - pass-through traffic*, 573–576
- virtualization requirements, 105–106
- deployment architecture, 2
- derived policies, 574
- description parameter (VB configuration), 485
- design solutions
 - application characteristics, 90–91
 - Application Optimizer (AO) requirements, 91–98
 - availability requirements, 99–100
 - management requirements, 100–103
 - network infrastructure, 82–90
 - performance and scalability metrics, 62–75
 - physical environment, 81–82
 - planning deployment, 77–78
 - platform requirements, 98–99
 - requirements collection and analysis, 78–79
 - scalability requirements, 99
 - security requirements, 103–105
 - site information, 80
 - Site Information Checklist, 82
 - user populations, 81
 - virtualization requirements, 105–106
- device CLI, acceleration monitoring, 453–459
- device configuration service (XML-API), 310–311
- device group context, 333
- device mode, 257
- device status service (XML-API), 312
- devices
 - alarms, 290–291
 - assigning application groups to, 352
 - boot sequence interruption, 256
 - context, 333
 - groups, assigning application groups to, 352
 - identity recovery, 269–270
 - locations, 271
 - management, 250
 - CLI (command-line interface)*, 260–261
 - CM (Central Manager)*, 261–266
 - CMS service*, 266–268
 - configuration of CM role*, 258
 - configuration of device mode*, 258
 - configuration of primary interface*, 257
 - configuration of settings*, 252–253, 256
 - groups*, 271–273
 - interfaces*, 250
 - performance and scalability metrics*, 73
 - registration*, 269–271
 - setup wizard*, 250, 252, 256–260

- managing reports, 295–296
- memory, performance and scalability metrics, 63–64
- monitoring charts, 291–295
- performance, 388–389, 392–393
- provisioned management, 273–274
 - integration with centralized authentication*, 278–289
 - RBAC*, 274–278
- status verification, 563
- DHCP (Dynamic Host Configuration Protocol), 250, 475
- diagnostic tests, 589–592
- Differentiated Services Code Point (DSCP), 358
- differentiated services code point. *See* DSCP, 54
- Digital Media Player (DMP), 415
- Digital Media System (DMS), 415
- digital signage, video acceleration, 415
- digital signatures, 93
- dir command, 300, 480
- Directed Mode, 149, 327–329
- Directed mode
 - configuration of WAN optimization, 338
 - UDP, 149
- directives, prepositioning, 448–449, 452
- Directory Replication Service. *See* DRS
- disabling
 - AOs, 416, 419–420, 423
 - configuration of WAN optimization, 331–333
 - EPM classification, 367
 - Telnet, 260
- discovery and analysis
 - application characteristics, 90–91
 - Application Optimizer (AO) requirements, 91
 - advanced features*, 92
 - CIFS AO*, 91–92
 - file services utilization*, 93
 - HTTP AO*, 95–96
 - MAPI AO*, 94–95
 - NFS AO*, 96
 - Replication Accelerator*, 98
 - SSL AO*, 97
 - Video AO*, 96–97
 - availability requirements, 99–100
 - deployment planning, 77–78
 - management requirements, 100–103
 - network infrastructure, 82
 - data center topology*, 86
 - Network Infrastructure Checklist*, 89–90
 - remote office topology*, 85–86
 - traffic flows*, 87–89
 - WAN topology*, 82, 84
 - platform requirements, 98–99
 - requirements collection, 78–79
 - scalability requirements, 99
 - security requirements, 103–105
 - site information, 80–82
 - virtualization requirements, 105–106
- disks
 - capacity
 - performance and scalability metrics*, 64–65
 - virtual blades*, 476
 - emulation parameter (VB configuration), 486
 - monitoring, CLI, 389
 - parameter (VB configuration), 486
- dispatchers, 366
- distortion, ART measurements, 396
- distribution switches (WANs)
 - WCCP enabled on, 217, 219–226
- DMP (Digital Media Player), 415
- DMS (Digital Media System), 415
- DNS (Domain Name System), 250, 475
- Do not set setting (policy map Accelerate parameter), 359
- domain component (RBAC), 274
- domain definitions, CM (Central Manager), 104
- Domain Name System (DNS), 250, 475
- domains, 274
 - assigning to users, 277
 - configuration, 276–277
 - Windows parameters, 281–283
- downgrades, software, 302–305
- DRE (Data Redundancy Elimination), 322–323
 - cold transfers, 564
 - monitoring in CLI, 392
- DRE compression history capacity, 65
- DRS (Directory Replication Service), 94
- DSCP (differentiated services code point), 54, 358
- DSCP Marking parameter (policy maps), 360
- dual data centers
 - asymmetric routing, 224–226
 - symmetric routing, 205
- Dynamic Host Configuration Protocol (DHCP), 250, 475
- dynamic limits, performance and scalability metrics, 63
 - device memory, 63–64
 - disk capacity, 64–65

- number of devices managed, 73
- peers and fan-out, 71, 73
- replication acceleration, 74–75
- TCP connections, 65–69
- virtual blades, 75
- WAN bandwidth and LAN throughput, 70–71

dynamic port assignment, 366–367

dynamic services, 121

dynamic shares, 92

E

e-mail compression, MAPI acceleration, 410

E1000 emulation method, 487

edge (WANs)

- data center traffic, 205
- WCCP enabled on routers, 212–217

edge target devices, repositioning, 447–449

Egress Methods for Intercepted Connections. *See* EMIC

egress traffic, EMIC, 145–149

Eject CD-ROM button (CM), 498

EMIC (Egress Methods for Intercepted Connections), 145–149

EMSMDB (Exchange Server STORE Electronic Messaging System Microsoft Data Base), 94

Enabled parameter (policy maps), 360

enabling

- AOs, 416, 419–420, 423
- CM services, 259
- configuration of WAN optimization, 331–333
- EPM classification, 367
- licenses, 330
- SSH, 260

encapsulation, directed mode, 328

encryption, Cisco Linux platform, 50–51

EndPoint Mapper (EPM) classification, 366–370

EndPoint Mapper. *See* EPM

Enterprise license, 62, 329

Entity type property (domain configuration page), 276

environment, requirements for deployment, 81–82

EPM (EndPoint Mapper), 94

- classification, 366–370
- setting (policy map Type parameter), 359

Error Message Book, 102, 290

error messages, monitoring virtual blades, 505–506

eviction, isolated, 323

Exchange Server STORE Electronic Messaging System Microsoft Data Base (EMSMDB), 94

EXEC mode (CLI commands), 595

executive reports, WAN optimization statistics, 393

eXtensible Markup Language (XML), 308

eXtensible Markup Language Application Programming Interface. *See* XML-API

F

failure detection, WCCP, 126–128

failure to boot, troubleshooting virtual blades, 506

fairness (TFO optimization), 322

fan-out, performance and scalability metrics, 71–73

fan-out ratio, 212

FCIP (Fibre Channel over IP), 98

File Replication Service (FRS), 95

file servers

- offloads, CIFS acceleration, 404
- requirements for deployment, 93

File Services Requirements Checklist, deployment requirements, 93

file write optimization, NFS, 409

Firewall Services Module (FWSM), 335

Firewall Switch Module (FWSM), 240

firewalls

- data centers, 240–246
 - FWSM connection display output*, 240–241
 - PIX/ASA configuration*, 243–246
 - server farm aggregation with FWSM*, 241–243
- TCP ports, 261
- TFO blacklist operation, 335
- troubleshooting, 584

floppy image parameter (VB configuration), 486

flow protection, WCCP, 128

flow types (traffic flow), 88

FlowAgent, 397

Fluke NetFlow Tracker, 381

Fluke Networks Visual Performance Manager, 309

forwarding methods, 123–125

FRS (File Replication Service), 95

FTP (File Transfer Protocol)

- baseline measurements, 551
- hot transfers, 565

FTP server

- backing up virtual blade files, 502
- restoring virtual blade files, 503

Full Optimization default policy action, 210
 Full Optimization setting (policy map Action parameter), 359
 FWSM (Firewall Services Module), 240, 335
 connection display output, 240–241
 server farm aggregation, 241–243

G

Gateway Load Balancing Protocol (GLBP), 83, 146
 Generic Routing Encapsulation (GRE), 490
 getBytesCount parameter (SSLStats service), 467
 getCIFSCliAvgThroughput parameters (CIFSSStats service), 463
 getCIFSCoreCount parameters (CIFSSStats service), 464
 getCIFSCoreEdgeTraffic parameters (CIFSSStats service), 464
 getCIFSEdgeCoreTraffic parameters (CIFSSStats service), 464
 getCIFSEdgeCount parameters (CIFSSStats service), 464
 getCM interface (device configuration service), 311
 getCMByName interface (device configuration service), 311
 getConnOptRate parameters (HttpStats service), 468
 getDeviceGroups interface (device configuration service), 311
 getDeviceStatus interface (device status service), 312
 getDiskCapacity parameters (CIFSSStats service), 464
 getDiskEncryptStatus interface (device status service), 312
 getDiskInformation interface (device status service), 312
 getDiskStatus interface (device status service), 312
 getErrorConnCount parameter (SSLStats service), 467
 getMaxConnReuseCount parameters (HttpStats service), 468
 getMonitoredApplications interface (Traffic Acceleration service), 394
 getOpenFileCount parameters (CIFSSStats service), 464
 getOptCIFSSessionCount parameters (CIFSSStats service), 464
 getOptConnCount parameter (SSLStats service), 467
 getOptConnCount parameters (HttpStats service), 468
 getRequestCount parameters (CIFSSStats service), 464
 getSessionCount parameters (MapiStats service), 469
 getSessionCount parameters (NfsStats service), 470
 getTotalConnCount parameter (SSLStats service), 467
 getTotalConnCount parameters (HttpStats service), 468
 getUnAccelConnCount parameter (SSLStats service), 467
 getUnAccelConnCount parameters (HttpStats service), 468
 getWAE interface (device configuration service), 311
 getWAEByName interface (device configuration service), 311
 getWAEs interface (device configuration service), 311
 getWAEsInGroup interface (device configuration service), 311
 getWAEsInGroupByName interface (device configuration service), 311
 getWANInfo interface (device configuration service), 311
 Gigabit Ethernet (GigE) interface, 256, 489
 GLBP (Gateway Load Balancing Protocol), 83, 146
 Global configuration mode (CLI commands), 597
 graceful shutdown, WCCP, 128
 GRE (Generic Routing Encapsulation), 490
 forwarding, 123–125
 groups
 devices, 271–273
 service
 placement, 130–131
 WCCP, 120–123
 guest OS boot image, installation of VB software, 482–483
 GUI (CM), 250
 acceleration monitoring, 460, 462
 alarms, 589–592
 configuring SSL acceleration services, 433–434, 437
 login page, 262

H

half-duplex, troubleshooting, 572–573
 hang conditions, troubleshooting virtual blades, 508–509
 hardware
 architecture, 49, 55
 appliances, 58–61
 licenses, 61–62
 router-integrated network modules, 56–58
 virtual blades, 476–477, 481–482
 WCCP supported, 136–137
 hash assignments, 125
 hash bucket distribution, WCCP scalability, 233–234
 hash function (WCCP), 233
 health monitor collect rate, LCM cycle, 266
 hostnames, 259
 Hot Standby Router Protocol (HSRP), 83
 hot transfers, 565

- HSRP (Hot Standby Router Protocol), 83
 - HTTP acceleration, 411–412
 - HTTP AO (HTTP Application Optimizer), 95–96
 - HTTP Requirements Checklist, 95–96
 - HTTP setting (policy map Accelerate parameter), 360
 - HttpStats service, 463, 467–468
-
- I**
 - ICMP (Internet Control Message Protocol), 347
 - IDE (Integrated Drive Electronics), 311, 486
 - IDS (Intrusion Detection System), 86
 - IDS/IPS (Intrusion Detection System/Intrusion Prevention System), 205
 - in-path deployment, 153–154
 - nonredundant branch offices, 154–158
 - redundant branch offices, 158161
 - serial inline clustering, 162–163
 - independent software vendors (ISV), 394
 - informs (SNMP), deployment requirements, 101
 - InfoVista Application Optimization Manager, 309
 - initial setup wizard, 250, 252, 256–260
 - inline interception, 139–143, 489–490
 - cabling guidelines, 143
 - InlineGroup configuration, 141
 - multiple routers, 140
 - one-armed routing, 141
 - operating modes, 140
 - troubleshooting, 577–579
 - inline vlan all command, 161
 - inline vlan number command, 579
 - inlineGroup
 - configuration, 155–161
 - management IP address, 157–158
 - inspect waas command, 584
 - inspect was command, 243
 - integrated development environment (IDE), 311
 - Integrated Drive Electronics (IDE), 486
 - Integrated Services Router. *See* ISR
 - IntegratedServicesEngineslot/port interface, 170
 - integration
 - centralized authentication (provisioned management)
 - RADIUS*, 288–289
 - TACACS+*, 286–287
 - Windows*, 280–286, 289
 - network, best practices, 150–151
 - WAN optimization and third-party systems, 393
 - ART monitoring*, 394–399
 - XML-API*, 394
 - Intel x86 hardware architecture, 49
 - interception
 - content switching, 143–145
 - inline, 139–143
 - cabling guidelines*, 143
 - InlineGroup configuration*, 141
 - multiple routers*, 140
 - one-armed routing*, 141
 - operating modes*, 140
 - policy-based routing, 139
 - policy-based routing interception, 196–199
 - redirection
 - failure detection*, 126, 128
 - flow protection*, 128
 - forwarding/return methods*, 123–125
 - graceful shutdown*, 128
 - load distribution*, 125–126
 - redirect lists*, 129
 - scalability*, 129
 - service group placement*, 130–131
 - troubleshooting, 577
 - inline interception*, 577–579
 - WCCP*, 579–583
 - WCCP, 119
 - hardware-based platforms*, 136–137
 - IOS Firewall*, 200–201
 - overview*, 120
 - service groups*, 120–123
 - WCCP configuration, 131–135
 - IOS*, 133–134
 - IOS with inbound redirection*, 134–135
 - router lists*, 132–133
 - interception method
 - Data Center, 533–534
 - Profile A, 513
 - Profile B, 520
 - Profile C, 525–526
 - Interface configuration mode (CLI commands), 597
 - interfaces
 - alarm status service, 312
 - bridging, VB resource configuration, 489
 - inline interception*, 489–490
 - WCCP interception*, 490–493

device configuration service, 311

device management, 250

- CLI (command-line interface)*, 260–261
- CM overview*, 261–266
- CMS service*, 266–268

device status service, 312

IntegratedServicesEngineslot/port, 170

managers, Cisco Linux platform, 51

network, NME-WAE family, 56

NME-WAE connectivity, 108

programmatic, 308–316

traffic acceleration service, 313

WAE

- configuring*, 108–111
- connectivity*, 107–111
- names*, 110
- standby interface feature*, 115–119

Internet Control Message Protocol (ICMP), 347

Internet service providers (ISP), 203

interruption, device boot sequence, 256

Intrusion Detection System (IDS), 86

Intrusion Detection System/Intrusion Prevention System (IDS/IPS), 205

Intrusion Prevention System (IPS), 86

IOS Firewall, 199–201, 335

IOS WCCP global configuration, 133–134

IOS WCCP inbound redirection configuration, 134–135

IP addresses

- allocation, 474
- management, 157–158
- multiple next-hop addresses, 198
- next-hop addresses, 197

IP forwarding, 145–146

ip inspect WAAS enable command, 584

ip name-server command, 283

ip wccp redirect exclude in command, 180

ip wccp redirect exclude in commands, 134

ip wccpservice groupaccelerated command, 581

IPS (Intrusion Prevention System), 86

ISO images, copying to virtual blade, 482

isolated eviction, DRE per-peer context architecture, 323

isolation, platforms and services, 475

ISP (Internet service providers), 203

ISR (Integrated Services Router), 56, 250, 252

ISV (independent software vendors), 394

J-K

jobs drawer, CM My WAN context homepage, 264

Kdump tool, 507

keepalives (TFO), 344

Kerberos authentication, 282

kernel crash dumping, 507

kernel-integrated virtualization services, 50

key management, CM, 261

L

L2 forwarding, 124–125

LACP (Link Aggregation Control Protocol), 111

LANs (local-area networks)

- configuration, Profile A, 517–519
- switches
 - large nonredundant deployment*, 176–177
 - WCCP configuration*, 178–181
 - WCCP interception*, 178
- throughput, performance and scalability metrics, 70–71

large nonredundant branch offices, off-path deployment, 174–181

large redundant branch offices, off-path deployment, 190–191, 194, 196

latency, DRE per-peer context architecture, 323

- RTT (round-trip time), 83

latency mitigation (TFO optimization), 321

Layer 2 bridging interface, 489

LCM (Local Central Manager) cycle, 51, 266

- data feed poll rate, 266
- health monitor collect rate, 266
- monitoring collect rate, 266

legacy CIFS, 406

legacy mode services, 262

LFNs (long fat networks), 339

licenses, 61–62

- activation, 478
- configuration, 330
- configuration of WAN optimization, 329–331
- enabling, 330
- virtual blades, 478

link aggregation, PortChannel, 111–115

Link Aggregation Control Protocol. *See* LACP

Linux Kernel Crash Dump (LKCD), 507

LKCD (Linux Kernal Crash Dump), 507
 load distribution, WCCP, 125–126
 local authentication, 280
 Local Central Management cycle. *See* LCM (Local Central Management) cycle
 local response handling
 CIFS acceleration, 405
 Windows print acceleration, 407
 local-area networks. *See* LANs
 logging mechanisms
 audit trails, 296–297
 SNMP configuration, 300–301
 system messages, 297–300
 login page (CM GUI), 262
 long fat networks (LFNs), 339
 loss mitigation (TFO optimization), 321
 low-compression ratios, troubleshooting, 584–587

M

mac-sticky command, 228
 managed service providers (MSP), 309
 management
 CM capabilities, 289
 alarms and monitors, 290–295
 backup and restoration of database, 305–307
 logging mechanisms, 296–301
 programmable interfaces, 308–316
 reports, 295–296
 software upgrades/downgrades, 302–305
 devices, 250
 CLI (*command-line interface*), 260–261
 CM (*Central Manager*), 261, 266
 CMS *service*, 266–268
 configuration of settings, 252–253, 256
 groups, 271–273
 interfaces, 250
 performance and scalability metrics, 73
 registration, 269–271
 setup wizard, 250, 252, 256–260
 protocols, 100
 requirements for deployment, 100
 CM and XML-API, 100
 Management Requirements Checklist, 103
 SNMP *community strings*, 101–102
 SNMP *traps/informs*, 101
 syslog servers, 102
 system, 250
 VBs (*virtual blades*), 476
 virtual blades, 500
 Management Informations Bases. *See* MIB
 management IP addresses, 157–158
 Management Requirements Checklist, 103
 manual configuration, device primary interface, 257
 map basic disable (#) command, 362
 map basic insert command, 362
 map basic list command, 362
 map basic move from (#) to (#) command, 362
 map command, 362
 MAPI (Messaging Application Programming Interface), 409–411
 MAPI AO (Messaging Application Programming Interface Application Optimizer), 94–95
 MAPI Requirements Checklist, 95
 MAPI setting (policy map Accelerate parameter), 360
 MapiStats service, 463, 468–469
 mask assignments, 126
 mask/value set distribution, WCCP scalability, 236–238
 match conditions, traffic classifiers, 355
 maximum segment size (MSS), 328
 medium nonredundant branch offices, off-path deployment, 163–167, 169–170
 NME-WAE, 170–171
 two-arm deployment, 171–174
 medium redundant branch offices, off-path deployment, 181–190
 memory
 dump files, locating, 507
 performance and scalability metrics, 63–64
 memory parameter (VB configuration), 486
 Messaging Application Programming Interface Application Optimizer. *See* MAPI AO
 Messaging Application Programming Interface. *See* MAPI
 metadata caching
 CIFS acceleration, 404
 Windows print acceleration, 407
 MIBs (Management Information Bases), 52, 101, 301, 381
 Microsoft Access-Based Enumeration. *See* ABE
 Microsoft Management Console (MMC), 500
 Microsoft Performance Monitor (Perfmon), 93
 Microsoft Remote Procedure Call. *See* MS-RPC connections
 Microsoft Volume Shadow Copy Services. *See* VSS
 Microsoft Windows Server Core 2008, 475

mitigation, 321

mkdir command, 300

MMC (Microsoft Management Console), 500

models

- appliances
 - WAE model 512, 60
 - WAE model 612, 60
 - WAE model 674, 61
 - WAE model 7341, 61
 - WAE model 7371, 61
 - WAVE model 274, 59
 - WAVE model 474, 59
 - WAVE model 574, 60
- router-integrated networks
 - NME-WAE model 302 (NME-WAE-302), 57
 - NME-WAE model 502 (NME-WAE-502), 57
 - NME-WAE model 522 (NME-WAE-522), 58

Monitor drawer

- CM My WAN context homepage, 264
- monitoring charts, 293–295

monitoring

- acceleration
 - CM GUI, 460, 462
 - device CLI, 453–459
 - XML-API, 463, 465–470
- charts, 291–295
- collect rate, LCM cycle, 266
- CPU utilization, 388–389
- DRE in CLI, 392
- facilities, Cisco Linux platform, 52
- virtual blades, 503–506
- WAN optimization, 370
 - automatic discovery statistics, 370–372
 - connection statistics, 373–374, 378, 380
 - integration with third-party visibility systems, 393–399
 - optimization statistics, 380–389, 392–393

MS Port Mapper setting (policy map Accelerate parameter), 359

MS-RPC (Microsoft Remote Procedure Call) connections, 94

MSP (managed service providers), 309

MSS (maximum segment size), 328

multi-data centers, 205

multiple next-hop addresses, 198

My WAN context homepage (CM), 262–265

N

NAM (Network Analysis Module), 106, 381, 397

Name Service Provider Interface. See NSPI

names, WAE interfaces, 110

NBAR (Network Based Application Recognition), 86, 381

negotiating policies (ATP), 365–366, 574

net use command, 550

NetFlow, 380

network interfaces (VBs), configuration, 487

NetQoS Performance Center, 309, 381

netstat command, 68

Network Analysis Module (NAM), 106, 381, 397

Network Based Application Recognition (NBAR), 381

Network File System Application Optimizer. See NFS AO

Network File System. See NFS

Network Infrastructure Checklist, 89–90

network interface card (NIC), 476

network module console, accessing, 250–252

Network Module Enhanced WAE. See NME-WAE

Network Operations Center. See NOC

network round trip time (NRTT), ART monitoring, 395

Network-Based Application Recognition. See NBAR

networks

- infrastructure, requirements for deployment, 82
 - data center topology, 86
 - Network Infrastructure Checklist, 89–90
 - remote office topology, 85–86
 - traffic flows, 87–89
 - WAN topology, 82–84
- Integration, best practices, 150–151
- interception
 - Cisco Linux platform, 52
 - content switching, 143–145
 - failure detection, 126–128
 - flow protection, 128
 - forwarding/return methods, 123–125
 - graceful shutdown, 128
 - inline, 139–143
 - load distribution, 125–126
 - policy-based routing, 139
 - redirect lists, 129
 - scalability, 129
 - service group placement, 130–131

- WCCP, 119–123
- WCCP configuration, 131–135
- WCCP hardware-phased platforms, 136–137
- interfaces, NME-WAE family, 56
- profiling, WAN optimization statistics, 380–385
- topologies
 - Data Center, 533
 - Profile A, 513
 - Profile B, 520
 - Profile C, 525
- next-hop addresses
 - multiple, 198
 - verifying, 197
- next-hop verify-availability route map command, 197
- NFS (Network File System), 408–409
- NFS AO (Network File System Application Optimizer), 96
- NFS Requirements Checklist, 96
- NFS setting (policy map Accelerate parameter), 360
- NfsStats service, 463, 470
- NIC (network interface card), 476
- NIC emulation parameter (VB configuration), 487
- NME (Network Module Enhanced) hardware, 56, 170
- NME-WAE, 108, 170–171
 - model 302 (NME-WAE-302), 57
 - model 502 (NME-WAE-502), 57
 - model 522 (NME-WAE-522), 58
- no normalization command, 230
- no vnc CLI command, 495
- no wccp ver 2 command, 128
- NOC (Network Operations Center), 274, 394
- nonredundant branch offices
 - in-path deployment, 154–158
 - off-path deployments
 - large offices, 174–181
 - NME-WAE, 170–171
 - small to medium office, 163–174
 - two-arm deployment, 171–174
- NRIT (network round trip time), ART monitoring, 395
- NSPI (Name Service Provider Interface), 95
- NTLM authentication, 282
- Number of Retransmits setting
 - ITACACS+ configuration, 286
 - RADIUS configuration, 288

O

- OAB (Outlook Address Book), 94
- object delivery acceleration, 410
- objects, read-aheads, 410
- OCSP (Online Certificate Status Protocol), 413
- off-path deployment, 163
 - IOS FW, 199–201
 - nonredundant branch offices
 - large offices, 174–181
 - NME-WAE, 170–171
 - small to medium office, 163–174
 - two-arm deployment, 171–174
 - policy-based routing interception, 196–199
 - redundant branch offices
 - large offices, 190–191, 194, 196
 - small to medium offices, 181–184, 186, 189–190
- one or more entities property (domain configuration page), 276
- Online Certificate Status Protocol (OCSP), 413
- Online state (device activation), 271
- operational latency, 323
- optimization
 - CM, 565–566
 - connections, 320–321
 - monitoring charts, 291–295
 - performance and scalability metrics
 - disk capacity, 64–65
 - number of devices managed, 73
 - peers and fan-out, 71–73
 - replication acceleration, 74–75
 - TCP connections, 65–69
 - virtual blades, 75
 - WAN bandwidth and LAN throughput, 70–71
 - WANs (wide-area networks), 319
 - ATP (Application Traffic Policy), 347–352, 355–370
 - automatic discovery, 324–327
 - configuration, 329–347
 - directed mode, 327–329
 - DRE (Data Redundancy Elimination), 322–323
 - monitoring and reporting, 370–389, 392–399
 - PLZ (Persistent LZ Compression), 324
 - TFO (Transport Flow Optimization), 320–322
- Optimization tab (monitoring charts), 292
- Optimize full default option, 364

organizational units (OUs), 282
 original connections, 320
 OUs (organizational units), 282
 Outlook Address Book (OAB), 94

P

packets, TCP reset (RST), 150
 Paessler Router Traffic Grapher (PRTG), 381
 PAgP (Cisco Port Aggregation Protocol), 111
 panic function, 507
 para-virtualization (PV) drivers, 487
 parameters
 CIFSStats service, 463–464
 defining policy maps, 358
 SSLStats service, 467
 Pass-through default option, 365
 pass-through traffic, troubleshooting, 573–576
 Passthrough setting (policy map Action parameter), 359
 payload aggregation, 410
 PBR (policy-based routing), 139, 227
 PCI (Peripheral Component Interconnect) expansion, 58
 peers, performance and scalability metrics, 71–73
 Pending state (device activation), 271
 per-peer context architecture, DRE, 323
 Perfmon (Microsoft Performance Monitor), 93
 performance, 3
 application
 CM optimization results, 565–566
 hot transfer measurements of FTP, 565
 measuring, 564
 testing tools, 566
 Windows File Transfer, 564–565
 baseline texts, 549
 FTP measurements, 551
 Windows file services, 549–551
 solution design, 62–63
 device memory, 63–64
 disk capacity, 64–65
 number of devices managed, 73
 peers and fan-out, 71, 73
 replication acceleration, 74–75
 TCP connections, 65–69
 virtual blades, 75
 WAN bandwidth and LAN throughput, 70–71

troubleshooting, 570
 application acceleration, 587–589
 firewall integration, 584
 half-duplex conditions, 572–573
 interception, 577–583
 low-compression ratios, 584–587
 pass-through traffic, 573–576
 performance improvement
 TFO capabilities, 321–322
 WAN optimization statistics, 386–388
 Performance Monitor (Perfmon), 93
 Peripheral Component Interconnect (PCI expansion), 58
 permissions, XML-API, 308
 Persistent LZ Compression. *See* Persistent LZ Compression
 physical environment, requirements for deployment, 81–82
 PIX (Private Internet eXchange), 335
 configuration, 243–246
 placement, service groups, 130–131
 planning deployment, 77–78
 application characteristics, 90–91
 Application Optimizer (AO) requirements, 91
 advanced features, 92
 CIFS AO, 91–92
 file services utilization, 93
 HTTP AO, 95–96
 MAPI AO, 94–95
 NFS AO, 96
 Replication Accelerator, 98
 SSL AO, 97
 Video AO, 96–97
 availability requirements, 99–100
 checklist, 78
 management requirements, 100
 CM and XML-API, 100
 Management Requirements Checklist, 103
 SNMP community strings, 101–102
 SNMP traps/informs, 101
 syslog servers, 102
 network infrastructure, 82
 data center topology, 86
 Network Infrastructure Checklist, 89–90
 remote office topology, 85–86
 traffic flows, 87–89
 WAN topology, 82–84
 platform requirements, 98–99

- requirements collection and analysis, 78–79
- scalability requirements, 99
- security requirements, 103–105
- site information, 80
 - physical environment*, 81–82
 - Site Information Checklist*, 82
 - types of sites*, 80
 - user populations*, 81
- virtualization requirements, 105–106
- Platform Requirements Checklist**, 98–99
- platform statistics, monitoring charts**, 291–295
- Platform tab (monitoring charts)**, 292
- platforms**
 - capacity, 65
 - isolation, 475
 - licenses, 61–62
 - requirements for deployment, 98–99
 - VBs (virtual blades), 477
 - virtual blades, 477
 - WCCP hardware-based, 136–137
- Plixer Scrutinizer**, 381
- PLZ (Persistent LZ Compression)**, 324
- point-to-multipoint tunnel interface configuration**, 148–149
- policies**
 - applied, 574
 - configurable elements, 53–54
 - configured, 574
 - derived, 574
 - negotiated, 574
 - peer's, 574
- policy maps (ATP)**, 358–365
- policy negotiation (ATP)**, 365–366
- policy prioritization**, 358
- policy-based routing (PBR)**, 227
- policy-based routing interception**, 196–199
- policy-engine application command**, 362
- populations, requirements for deployment**, 81
- port mappers**, 366–367
- PortChannel**
 - configuring, 112–115
 - link aggregation, 111
- PortFast feature**, 155, 160
- ports**
 - CFIS, 262
 - TCP, 261
 - WAFS, 262
- Position parameter (policy maps)**, 360
- predictor method (ACE)**, 239–240
- prepositions**, 404
 - configuring, 447–452
 - core locations, 447–449
 - directives, 448–449, 452
 - edge target devices, 447–449
- primary servers, configuration**, 286
- print services capacity**, 65
- prioritization, policies**, 358
- Private Internet eXchange (PIX)**, 335
- Profile A**, 512
 - interception method, 513
 - LAN switch configuration, 517–519
 - network topology, 513
 - requirements, 513
 - WAE configuration, 513–515
 - WAE placement, 513
 - WAN router configuration, 516–517
- Profile B**, 519
 - interception method, 520
 - network topology, 520
 - requirements, 519
 - WAE configuration, 520–522
 - WAE placement, 520
 - WAN router configuration, 522–532
- Profile C**, 524
 - interception method, 525–526
 - network topology, 525
 - requirements, 524
 - WAE configuration, 526–528
 - WAE placement, 525–526
- programmatic interfaces**, 308–309
 - accessible data, 310–316
 - vendor support, 309–310
- Property Configuration page (CM GUI page)**, 267
- protocols**
 - management, 100
 - WCCP, 119
 - configuring*, 131–135
 - failure detection*, 126–128
 - flow protection*, 128
 - forwarding/return methods*, 123–125
 - graceful shutdown*, 128
 - hardware-based platforms*, 136–137
 - load distribution*, 125–126

- overview*, 120
- redirect lists*, 129
- scalability*, 129
- service group placement*, 130–131
- service groups*, 120–123

provisioned management, 273–274

- integration with centralized authentication, 278–280
 - RADIUS*, 288–289
 - TACACS+*, 286–287
 - Windows*, 280–286, 289

- RBAC, 274–278

proxy (TFO), interaction with optimized connections, 321

PRTG (Paessler Router Traffic Grapher), 381

PV (para-virtualization) drivers, 487

Q

quickstart guide

- baseline performance tests, 549
 - FTP measurements*, 551
 - Windows file services*, 549–551
- configuring client-side WAE with Setup Wizard, 559–563
- configuring server-side WAAC device with WCCPv2, 555–559
- configuring WAAS CM with Setup Wizard, 552–555
- setup verification, 563
- test lab setup, 547–548

R

RADIUS (Remote Authentication Dial In User Service), 103, 288–289

RADIUS Servers and Ports setting (RADIUS configuration), 289

RAID (Redundant Array of Inexpensive Disks), 56

RBAC (Role-Based Access Control), 103, 249, 273–278

read requests, CIFS acceleration, 404

read-aheads

- CIFS acceleration, 404
- object, MAPI acceleration, 410

Read-Only Active Directory Services (ROADS), 475

recovery, device identity, 269–270

recovery point objective (RPO), 98

recovery time objective (RTO), 98

redirect lists, WCCP, 129

redirecting traffic, 123–125

redirection, WCCP

- failure detection, 126–128
- flow protection, 128
- forwarding/return methods, 123–125
- graceful shutdown, 128
- load distribution, 125–126
- redirect lists, 129
- scalability, 129
- service group placement, 130–131

Redundant Array of Inexpensive Disks. *See* RAID

redundant branch offices

- in-path deployment, 158–161
- off-path deployments
 - large offices*, 190–191, 194, 196
 - small to medium offices*, 181–186, 189–190

registration

- CM, verification, 563
- CMS status, 268
- devices, 269–271

remediation actions, alarms, 291

Remote Authentication Dial In User Service (RADIUS), 103

remote office topology, requirements for deployment, 85–86

Remote Procedure Call. *See* MS-RPC connections

remote sites

- Profile A, 512
 - interception method*, 513
 - LAN switch configuration*, 517–519
 - network topology*, 513
 - requirements*, 513
 - WAE configuration*, 513–515
 - WAE placement*, 513
 - WAN router configuration*, 516–517

- Profile B, 519
 - interception method*, 520
 - network topology*, 520
 - requirements*, 519
 - WAE configuration*, 520–522
 - WAE placement*, 520
 - WAN router configuration*, 522–532

- Profile C, 524
 - interception method*, 525–526
 - network topology*, 525
 - requirements*, 524
 - WAE configuration*, 526–528
 - WAE placement*, 525–526

removing

- application groups, 349
- traffic classifiers, 353

replication acceleration

- configuration of WAN optimization, 345–347
- performance and scalability metrics, 74–75
- requirements for deployment, 98

report drawer, CM My WAN context homepage, 264**reporting, WAN optimization, 370**

- automatic discovery statistics, 370–372
- connection statistics, 373–374, 378–380
- integration with third-party visibility systems, 393–399
- management, 295–296
- optimization statistics, 380–389, 392–393

resources, virtual blades

- configuration, 484–493
- hardware, 481–482

restoration

- CM database, 305–307
- virtual blades, 501–502

retransmission delay, ART monitoring, 395**retrieveAlarmByName interface (alarm status service), 312****retrieveAlarmBySeverity interface (alarm status service), 312****retrieveAllAlarms interface (alarm status service), 312****retrieveAppTrafficStats interface (Traffic Acceleration service), 394****retrieveCacheObjectCount parameters (CIFStats service), 464****retrieveCacheUtilization parameters (CIFStats service), 464****retrieveClientConnCount parameters (MapiStats service), 469****retrieveConnection interface (Traffic Acceleration service), 394****retrieveCPUUtilization interface (traffic acceleration service), 313, 394****retrieveCurrentStats parameter (VideoStats service), 467****retrieveDataReadStats parameters (MapiStats service), 469****retrieveHistoricalStats parameter (VideoStats service), 467****retrieveNFSTypeStats parameters (NfsStats service), 470****retrieveRequestHitRate parameters (CIFStats service), 464****retrieveRequestTypeStats parameters (MapiStats service), 469****retrieveRequestTypeStats parameters (NfsStats service), 470****retrieveResponseStats parameters (MapiStats service), 469****retrieveResponseStats parameters (NfsStats service), 470****Return Material Authorization (RMA), 502****return methods, 123–125****reverse Telnet sessions, accessing VB console, 496****RMA (Return Material Authorization), 502****ROADS (Read-Only Active Directory Services), 475****role component (RBAC), 274****role definitions, CM (Central Manager), 104****Role-Based Access Control (RBAC), 103****role-based access control. See RBAC, 249, 273****roles**

- assigning to users, 277

- configuration, 275–276

round-trip time (RTT) latency, 83**routed mode (ACE), 231–232****routed mode (ACE), 227, 230–232****router-integrated network modules, 56–57**

- hardware architecture, 49

- NME-WAE model 302 (NME-WAE-302), 57

- NME-WAE model 502 (NME-WAE-502), 57

- NME-WAE model 522 (NME-WAE-522), 58

routers

- data center traffic, 203

- lists, WCCP, 132–133

- WAN

large redundant off-path deployment, 191, 194

small to medium redundant deployment, 186

routing

- asymmetric, dual data centers, 224–226

- PBR (policy-based routing), 227

routing traffic

- asymmetric flow, 207–208

- symmetric flow, 205

RPO (recovery point objective), 98**RST (TCP reset), 150****RTL8139 emulation method, 487****RTO (recovery time objective), 98****RTT (round-trip time) latency, 83****running state (VBs), 493**

S

safe data caching, CIFS acceleration, 404

scalability

bandwidth (TFO optimization), 322

CM, 262

compression, 323

metrics, solution design, 62–63

device memory, 63–64

disk capacity, 64–65

number of devices managed, 73

peers and fan-out, 71–73

replication acceleration, 74–75

TCP connections, 65–69

virtual blades, 75

WAN bandwidth and LAN throughput, 70–71

requirements for deployment, 99

transparent interception, 233

ACE, 239–240

WCCP, 233–239

WCCP, 129

WAN optimization statistics, 388–393

SCCM (System Center Configuration manager), 500

schedules, prepositioning, 449, 452

SCOM (System Center Operations Manager), 500

scripts, 250

secondary servers, configuration, 286

Secure Sockets Layer (SSL) encrypted session, 51

security

CM, 261

deployment requirements, 105

requirements for deployment, 103–105

Security Word setting (ITACACS+ configuration), 287

serial connection, 250

serial inline clustering, in-path deployment, 162–163

Server Core 2008, 475

server farm aggregation, 241–243

server farm distribution

multiple WAAS clusters, 211–212

WAAS placement, 209, 211

server load balancing, 227

ACE bridged mode deployment, 227–230

ACE routed mode deployment, 230–232

server offloads

CIFS acceleration, 404

video acceleration, 41

server response time, ART monitoring, 395

Server Virtualization Validation Program (SVVP), 475

server-side WAAS device, 555–559

service level agreements (SLA), 393

service-module ip addressaddrmask command, 170

service-module ip default-gatewayaddr command, 171

services

groups

placement, 130–131

WCCP, 120–123

isolation, 475

XML-API AOs, 463

CIFSStats, 463–465

HttpStats, 467–468

HttpStats service, 463

MapiStats, 468–469

MapiStats service, 463

NfsStats, 470

NfsStats service, 463

SSLStats, 466–467

SSLStats service, 463

VideoStats, 467

VideoStats service, 463

set ip next-hop command, 139

setip next-hop command, 198

setup command, 253

setup time, ART monitoring, 395

setup wizard, 250–260

Setup Wizard

configuring client-side WAE, 559–563

configuring WAAS CM, 552–555

Shared Encryption Key setting (RADIUS configuration), 289

show accelerator command, 419, 587

show auto-discovery blacklist command, 372

show auto-discovery list command, 372

show clock command, 283

show cms info command, 267, 563

show command, 343

show conn longx command, 240

show disk details CLI command, 480

show disks details command, 591

show hardware command, 577

show interface command, 110–111, 149, 256, 572

show interface inlinegroup slot/number command, 578

show interface PortChannel command, 113–115

show interface Standby command, 117

- show ip wccp command, 580
- show ip wccpservice groupdetail command, 579
- show license command, 478
- show running-config command, 336, 577
- show statistics accelerator CIFS details command, 588
- show statistics application command, 383
- show statistics auto-discovery blacklist command, 371
- show statistics auto-discovery command, 370
- show statistics connection command, 373–375, 454, 588
- show statistics connection conn-idid command, 585
- show statistics connection optimized command, 588
- show statistics connection pass-through command, 574
- show statistics dre command, 392
- show statistics pass-through command, 575
- show tcam interface command, 581
- show tech-support command, 593
- show virtual-blade command, 481, 503
- show virtual-bladenumblockio command, 504
- show virtual-bladenuminterface command, 504
- show wccp gre command, 583
- show wccp routers command, 582
- show wccp services command, 582
- show wccp status command, 582
- showstatistics acceleratoracceleratordetail command, 457
- Simple Network Management Protocol. *See* SNMP
- Simple Object Access Protocol (SOAP), 308
- simplicity, 262
- site information, requirements for deployment, 80–82
- Site Model 2 (remote office topology), 86
- SLAs (service level agreements), 393
- slow-start mitigation (TFO optimization), 321
- small nonredundant branch offices, off-path deployment, 163–170
 - NME-WAE, 170–171
 - two-arm deployment, 171–174
- small redundant branch offices, off-path deployment, 181–186, 189–190
- SNMP (Simple Network Management Protocol), 52, 249
 - Alarm Book, 102
 - community strings, 101–102
 - configuration, 300–301
 - Management Information Bases, 101
 - MIBs, 381
 - traps/informs, 101
- snmp-server command, 300
- SOAP (Simple Object Access Protocol), 308
- soapUI tool, accessing XML-API data, 313–316
- software
 - architecture, 50–55
 - file entry, 302–303
 - installation on VB, 482–483
 - upgrades/downgrades, 302–305
- Solarwinds Orion Family, 381
- solution architecture, 2
- solution design
 - application characteristics, 90–91
 - Application Optimizer (AO) requirements, 91
 - advanced features*, 92
 - CIFS AO*, 91–92
 - file services utilization*, 93
 - HTTP AO*, 95–96
 - MAPI AO*, 94–95
 - NFS AO*, 96
 - Replication Accelerator*, 98
 - SSL AO*, 97
 - Video AO*, 96–97
 - availability requirements, 99–100
 - data center network integration, 212–246
 - management requirements, 100–103
 - network infrastructure, 82
 - data center topology*, 86
 - Network Infrastructure Checklist*, 89–90
 - remote office topology*, 85–86
 - traffic flows*, 87–89
 - WAN topology*, 82–84
 - performance and scalability metrics, 62–75
 - planning deployment, 77–78
 - platform requirements, 98–99
 - requirements collection and analysis, 78–79
 - scalability requirements, 99
 - security requirements, 103–105
 - site information, 80–82
 - virtualization requirements, 105–106
- spoof-client-ip feature, 123
- SSH, enabling, 260
- ssh-key-generate command, 260
- SSL
 - acceleration, 412–414, 432–447
 - acceleration charts, 566
 - AO (SSL Application Optimizer), 97
- SSL Requirements Checklist, 97
- SSLStats service, 463, 466–467
- standby command, 109

standby interface feature, 115–119

static limits, performance and scalability metrics, 63

- device memory, 63–64
- disk capacity, 64–65
- number of devices managed, 73
- peers and fan-out, 71, 73
- replication acceleration, 74–75
- TCP connections, 65–69
- virtual blades, 75
- WAN bandwidth and LAN throughput, 70–71

static TCP buffering, configuring WAN optimization, 339–345

statistics

- automatic discovery, 370–372
- connection, 373–374, 378–380
- WAN optimization, 380
 - device and system performance*, 388–389, 392–393
 - network profiling*, 380–385
 - performance improvement*, 386–388

status (device), verification, 563

stopped state (VBs), 493

stopping virtual blades, 496–497

stream splitting, 415

SuperAgent Aggregator, 84

SVI (switched virtual interface), 178

SVVP (Server Virtualization Validation Program), 475

switched virtual interface (SVI), 178

switches, LAN

- large nonredundant deployment, 176–177
- WCCP configuration, 178–181
- WCCP interception, 178

symmetric routing, dual data centers, 205

SYN (synchronize) packets, TCP, 325–326

SYN/ACK (synchronize and acknowledge) packets, TCP, 325–326

synchronize. *See* SYN (synchronize) packets

syslog messages, VB faulty shutdown, 505

syslog servers, deployment requirements, 102

sysreport, 592–593

systems

- management, 250
- performance, 388–389, 392–393
- provisioned management, 273–274
 - integration with centralized authentication*, 278–289
 - RBAC*, 274–278
- timers (CM), 267

System Center Configuration manager (SCCM), 500

System Center Operations Manager (SCOM), 500

system dashboard charts (admin drawer), 264

system messages, 297–300

system reports, 592–593

T

TAC (Technical Assistance Center), 55, 475, 509, 592

TACACS (Terminal Access Controller Access-Control System), 103

TACACS+

- authentication, 286–287
- configuration, 286–287
- servers, 287

TCO (total cost of ownership), 474

TCP (Transmission Control Protocol), 319, 325–326

- buffering, configuration of WAN optimization, 339–345
- connections, performance and scalability metrics, 65–69
- ports, 261

TCP reset (RST) packets, 150

Technical Assistance Center (TAC), 55, 475, 509, 592

Telnet, disabling, 260

Terminal Access Controller Access-Control System (TACACS), 103

terminal emulation software, 250

test lab setup (quickstart guide), 547–548

test self-diagnosticstest command, 590

tests

- baseline performance, 549
 - FTP measurements*, 551
 - Windows file services*, 549–551
- diagnostic tests, 589–592
- lab setup for quickstart guide, 547–548
- tools, 566

TFO (Transport Flow Optimization), 320

- automatic discovery (TFO AD), 324–326
 - blacklist operation*, 333–337
 - directed mode*, 328
- blacklist operation, configuration, 336
- buffer settings, configuration, 342–344
- performance improvements, 321–322
- proxy interaction with optimized connections, 321

TFO AD (TFO automatic discovery), 324–326

- blacklist operation, 333–337
- directed mode, 328

- TFO Only setting (policy map Action parameter), 359
- TFO with Data Redundancy Elimination setting (policy map Action parameter), 359
- TFO with LZ Compression setting (policy map Action parameter), 359
- third-party visibility systems (WAN optimization), 393
 - ART monitoring, 394–399
 - XML-API, 394
- throughput, LANs, 70–71
- Time to Wait setting (ITACACS+ configuration), 286
- Time to Wait setting (RADIUS configuration), 288
- top-level context, 263
- topologies
 - Data Center, 203, 533
 - asymmetric traffic flows*, 207–208
 - dual data centers*, 205
 - multi-data centers*, 205
 - requirements for deployment*, 86
 - server farm distribution*, 209–212
 - single WCCP service groups*, 209
 - WAAS placement*, 205
 - WAN edge*, 205
 - Profile A, 513
 - Profile B, 520
 - Profile C, 525
 - remote office, 85–86
 - WANs, 82–84
- total cost of ownership (TCO), 474
- total transaction time, ART monitoring, 395
- traffic
 - egress, 145–149
 - flows, requirements for deployment, 87–89
 - forwarding/return methods, 123–125
 - nonredundant in-path branch office topology, 155–157
 - redundant in-path branch office topology, 158–161
 - two-arm deployment, 172–173
- traffic acceleration service (XML-API), 312–313
- traffic classifiers (ATP), 352, 355–357
- Traffic tab (monitoring charts), 292
- transaction time, ART monitoring, 395
- Transmission Control Protocol. *See* TCP
- Transparent GRE packets received counter, 583
- transparent integration, MAPI acceleration, 410
- transparent interception methods, 212
 - scalability, 233
 - ACE*, 239–240
 - WCCP*, 233–239

- server load balancing, 227
 - ACE bridged mode deployment*, 227–230
 - ACE routed mode deployment*, 230–232
- WCCP, 212–213, 216
 - WAN distribution switches*, 217–226
 - WAN edge routers*, 212–217
- Transparent non-GRE non-WCCP packets received counter, 583
- Transparent non-GRE packets received counter, 583
- Transport Flow Optimization. *See* TFO
- Transport license, 62, 329
- traps (SNMP), 101
- troubleshooting
 - diagnostic tests, 589–592
 - system reports, 592–593
 - VBs (virtual blades), 506
 - BSOD (Blue Screen of Death)*, 507
 - failure to boot*, 506
 - hang conditions*, 508–509
 - WAAS deployment, 570–571
 - application acceleration*, 587–589
 - firewall integration*, 584
 - half-duplex conditions*, 572–573
 - interception*, 577–583
 - low-compression ratios*, 584–587
 - pass-through traffic*, 573–576
- tunnel interface configuration, 147–149
- two-arm deployment, 171–174
- type command, 300
- Type parameter (policy maps), 358
- type-tail /local1/errorlog/virtual-blade command, 506
- type-tail command, 300

U

- UDP (User Datagram Protocol), 149, 320
- universally-unique identifiers (UUID), 54
- upgrades, software, 302–305
- Use CD-ROM button (CM), 498
- user authentication, 278–289, 474
- user component (RBAC), 274
- User Datagram Protocol (UDP), 149, 320
- user groups, 279
 - Active Directory, 280
 - configuration, 281

users

- assigning roles and domains, 277
- nonlocal, authentication, 278–279
- populations, requirements for deployment, 81

UUIDs (universally-unique identifiers), 54

V

Validate Software File Settings, 303

values (MSS), 344

VBs (virtual blades), 473–476

- accessing console, 495–496
- capacity, 477
- changing boot sequence, 497–500
- creating, 478–482
 - activation of license, 478*
 - enabling virtualization, 479–480*
 - guest OS boot image, 482–483*
 - hardware resources, 481–482*
 - resource configuration, 484–493*
 - verification of license, 478*

hardware emulation, 476–477

management, 476, 500–502

monitoring, 503–506

platforms, 477

starting, 493–494

stopping, 496–497

troubleshooting, 506–509

/vbspace disk partition, 480–481

vendors, XML-API integration, 309–310

verification

- next-hop addresses, 197
- setup verification, 563

video acceleration, 414–415, 423–425

Video AO (Video Application Optimizer), 96–97

Video license, 62, 330

Video Requirements Checklist, requirements for deployment, 96–97

Video setting (policy map Accelerate parameter), 360

video-on-demand caching, 415

VideoStats service, 463, 467

viewing

- traffic classifiers, 353
- CMS registration status, 268

VIP (virtual IP) addresses, 146

VirtIO, 486–487

Virtual Blade configuration pane, 484

Virtual Blade Entries window, 484

Virtual Blade license, 62

virtual blades. *See* VBs

virtual IP (VIP) addresses, 146

virtual local area networks. *See* VLANs

Virtual Network Computing (VNC), 495

Virtual Private Network (VPN), 345

Virtual Router Redundancy Protocol (VRRP), 83

Virtual-Blade license, 330

virtual-blade n start command, 494

virtual-bladenumbrcd eject command, 498

virtual-bladenumbrkill-save-state command, 506

virtual-bladenumbrsave command, 502

virtualization, 473

- overview, 473–475
- requirements for deployment, 105–106
- virtual blades (VB), 475–476
 - accessing console, 495–496*
 - changing boot sequence, 497–500*
 - creating virtual blades, 478–493*
 - hardware emulation, 476–477*
 - management, 476, 500–502*
 - monitoring, 503–506*
 - platforms and capacity, 477*
 - starting virtual blades, 493–494*
 - stopping, 496–497*
 - troubleshooting, 506–509*

Virtualization Requirements Checklist, deployment requirements, 106

virtualization services, 50

visibility systems, third-party, 393–399

VLANs (virtual local area network)

- nonredundant in-path branch office topology, 155–157
- redundant in-path branch office topology, 158–161

VNC, accessing VB console, 496

VoIP VLANs (voice over IP virtual area networks)

- nonredundant in-path branch office topology, 155–157
- redundant in-path branch office topology, 158–161

Volume Shadow Copy Services (VSS), 92

VPNs (Virtual Private Network), 345

VRRP (ViSee GLBPrtual Router Redundancy Protocol), 83

VSS (Microsoft Volume Shadow Copy Services), 92

W

WAAS Mobile, 3

WAE (Wide-Area Application Engine), 49

client-side, 559–563

configuration

Data Center, 534, 536

large nonredundant deployments, 177–178

large redundant off-path deployment, 194

Profile A, 513, 515

Profile B, 520, 522

Profile C, 526, 528

WCCP in WAN edge routers, 216–217

content switching, 143–145

in-path deployments, 153–154

nonredundant branch offices, 154–158

redundant branch offices, 158–161

serial inline clustering, 162–163

inline interception, 139–143

cabling guidelines, 143

InlineGroup configuration, 141

multiple routers, 140

one-armed routing, 141

operating modes, 140

interfaces

configuring, 108–111

connectivity, 107–111

names, 110

standby interface feature, 115–119

off-path deployments, 163

IOS FW, 199–201

nonredundant branch offices, 163–181

policy-based routing interception, 196–199

redundant branch offices, 181–196

models, 60–61

PBR, 139

placement

Data Center, 533–534

Profile A, 513

Profile B, 520

Profile C, 525–526

PortChannel, 111–115

subnet WAE configuration, 189

WAFS (Wide Area File Services), 262

WAFS Transport setting (policy map Type parameter), 359

WANs (wide-area networks), 1

bandwidth, performance and scalability metrics, 70–71

configuration

Data Center, 537, 540, 543

Profile A, 516–517

Profile B, 522, 524, 528–532

distribution switches, WCCP enabled on, 217, 219–224, 226

edge

data center traffic, 205

WCCP enabled on routers, 212–217

nonredundant in-path branch office topology, 155–157

optimization capabilities, 319

ATP (Application Traffic Policy), 347–370

automatic discovery, 324–327

configuration, 329–347

directed mode, 327–329

DRE (Data Redundancy Elimination), 322–323

monitoring and reporting, 370–374, 378–399

PLZ (Persistent LZ Compression), 324

TFO (Transport Flow Optimization), 320–322

redundant in-path branch office topology, 158–161

router configuration

large redundant off-path deployment, 191, 194

small to medium redundant deployment, 186

topology

case study, 511–512

requirements for deployment, 82, 84

WAVE (Wide Area Virtualization Engine appliance), 473

branch office virtualization, 473

accessing VB console, 495–496

changing VB boot sequence, 497–500

creating virtual blades, 478–493

hardware emulation, 476–477

monitoring virtual blades, 503–506

overview, 473–475

platforms and capacity, 477

starting virtual blades, 493–494

stopping virtual blades, 496–497

troubleshooting virtual blades, 506–509

VB management, 500–502

virtual blades, 475–476

devices (Wide-Area Virtualization Engine), 49

models, 59–60

WCCP (Web Cache Communication Protocol), 78, 119, 212–213, 216
 configuring, 131–135
 hardware-based platforms, 136–137
 interception, 490–493, 579–583
 IOS Firewall, 200–201
 on LAN switch, 178
LAN switch configuration, 178–181
 overview, 120
 redirection
 failure detection, 126, 128
 flow protection, 128
 forwarding/return methods, 123–125
 graceful shutdown, 128
 load distribution, 125–126
 redirect lists, 129
 scalability, 129
 service group placement, 130–131
 scalability, 233–238
 client distribution, 234–235, 238–239
 hash bucket distribution, 233–234
 hash function, 233
 mask/value set distribution, 236–238
 service groups, 120–123
 WAN distribution switches, 217–226
 WAN edge routers, 212–217
wccp tcp-promiscuous service group command, 170
WCCPv2 (Web Cache Coordination Protocol version 2), 52, 120, 490, 555–559
Web Services Definition Language (WSDL), 310
web-cache service, 120
well-known services, 120
Wide Area File Services (WAFS), 262
Wide-Area Application Engine. *See* **WAE**
windows
 Add/Edit Interface, 487
 Virtual Blade Entries, 484
Windows
 domain parameters, 281–283
 locating memory dump files, 507
Windows Authentication, 103, 280–286, 289
Windows file services, baseline measurements, 549–551
Windows File Transfer, performance measurements, 564–565
Windows Media Technologies (WMT), 415
Windows on WAAS (WoW), 106, 475
Windows print acceleration, 407–408

WMT (Windows Media Technologies), 415
WOC (WAN Optimization Controller) Distortion, 396
WoW (Windows on WAAS), 106, 475
write memory command, 253
write-behinds, 404
WSDL (Web Services Definition Language), 310

X-Z

XML (eXtensible Markup Language), 308
XML-API (eXtensible Markup Language Application Programming Interface), 249
 acceleration monitoring, 463
 CIFSStats service, 463, 465
 HttpStats service, 463, 467–468
 MapiStats service, 463, 468–469
 NfsStats service, 463, 470
 SSLStats service, 463, 466–467
 VideoStats service, 463, 467
 CM capabilities, 308–309
 accessible data, 310–313
 soapUI tool, 313–316
 vendor support, 309–310
 deployment requirements, 100
 permissions, 308