

Planning Your WAAS Network

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

**Note**

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

This chapter contains the following sections:

- [Checklist for Planning Your WAAS Network, page 2-1](#)
- [Site and Network Planning, page 2-4](#)
- [About Autoregistration and WAEs, page 2-8](#)
- [Identifying and Resolving Interoperability Issues, page 2-10](#)
- [WAAS Devices and Device Mode, page 2-16](#)
- [Calculating the Number of WAAS Devices Needed, page 2-17](#)
- [Supported Methods of Traffic Redirection, page 2-18](#)
- [Access Lists on Routers and WAEs, page 2-24](#)
- [WAAS Login Authentication and Authorization, page 2-25](#)
- [Logically Grouping Your WAEs, page 2-26](#)
- [Data Migration Process, page 2-28](#)

Checklist for Planning Your WAAS Network

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

WAAS software supports the following three typical network topologies:

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

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

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

Planning Checklist

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

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



Note

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

To plan your network, follow these guidelines:

1. Complete the sizing phase that includes the following tasks:
 - Determine which locations in your existing network require WAAS optimization (for example, which branch offices and data centers).
 - Determine the number and models of the WAAS devices that are required for each location. Some key factors in this selection process is the WAN bandwidth, the number of users, and the expected use. Various hardware configurations are possible (for example, different hard disk models and RAM size). Consider running a cluster of WAEs where additional scalability and or failover is required. For more information, see the [“Calculating the Number of WAAS Devices Needed” section on page 2-17.](#)
 - Verify that you have purchased sufficient licenses to cover your needs.
2. Plan for management as follows:
 - Complete site and network planning (for example, obtain the IP and routing information including IP addresses and subnets, routers and default gateway IP addresses, and the hostnames for the devices). See the “Checklist of WAAS Network System Parameters” table in the *Cisco Wide Area Application Services Quick Configuration Guide*.
 - Determine the login authentication and login authorization methods (for example, external RADIUS, TACACS+, Windows domain servers) and accounting policies that you want your WAAS Central Managers and WAEs to use. For more information, see [Chapter 6, “Configuring Administrative Login Authentication, Authorization, and Accounting.”](#)
 - For security purposes, plan to change the predefined password for the predefined superuser account immediately after you have completed the initial configuration of a WAE. For more information, see [“WAAS Login Authentication and Authorization” section on page 2-25.](#)

- Determine if you need to create any additional administrative accounts for a WAAS device. For more information, see [Chapter 7, “Creating and Managing Administrator User Accounts and Groups.”](#)
 - Determine if you should group your WAEs into logical groups. For more information, see the [“Logically Grouping Your WAEs” section on page 2-26.](#)
 - Determine which management access method to use. By default, Telnet is used but SSH may be the preferred method in certain deployments. For more information, see the [“Configuring Login Access Control Settings for WAAS Devices” section on page 6-7.](#)
3. Plan for application optimization as follows:
- Determine and resolve router interoperability issues (for example, the supported hardware and software versions, router performance with interception enabled). For more information, see the [“Site and Network Planning” section on page 2-4.](#)
 - Determine the appropriate interception location when the data center or branch office is complex (for example, if your existing network uses a hierarchical topology).
 - Determine which WAAS services to deploy (for example, Wide Area File Services [WAFS] services, WAAS print services, and WAAS application acceleration). For more information about the different WAAS services, see [Chapter 1, “Introduction to Cisco WAAS.”](#)
 - Determine which WAAS software licenses to install. Software licenses enable specific WAAS services. For more information about installing software licenses, see the [“Managing Software Licenses” section on page 9-3.](#)
 - Determine which traffic interception methods to use in your WAAS network (for example, inline mode; WCCP Version 2 or policy-based routing (PBR) for promiscuous mode; DFS or NetBIOS for WAFS-only traffic). For more information, see the [“Supported Methods of Traffic Redirection” section on page 2-18.](#)



Note WCCP works only with IPv4 networks.

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



Note IP ACLs that are defined on a router take precedence over the IP ACLs that are defined on the WAE. For more information, see the [“Access Lists on Routers and WAEs” section on page 2-24.](#)

- Determine whether you need to define IP ACLs on the WAEs. For more information, see the [“Access Lists on Routers and WAEs” section on page 2-24.](#)



Note IP ACLs that are defined on a WAE take precedence over the WAAS application definition policies that are defined on the WAE. For more information, see the [“Precedence of IP ACLs and Application Definition Policies on WAEs” section on page 8-3.](#)

- If PBR is to be used, determine which PBR method to use to verify PBR next-hop availability for your WAEs. For more information, see the [“Methods of Verifying PBR Next-Hop Availability” section on page 4-38.](#)
- If you plan to deploy WAFS services, determine whether transparent or nontransparent interception methods (DFS or NetBIOS) should be used to intercept and redirect WAFS traffic to the local WAE. For more information, see the [“Request Redirection of CIFS Client Requests” section on page 4-48.](#)

- Determine the major applications for your WAAS network. Verify whether the predefined application definition policies cover these applications and whether you should add policies if your applications are not covered by these predefined policies. For a list of the predefined application definition policies, see [Appendix A, “Default Application Policies.”](#)
- Determine the print services configuration. For more information, see [Chapter 13, “Configuring and Managing WAAS Legacy Print Services.”](#)
- Consider day zero migration of file systems if file servers are to be centralized in the process. For more information, see the [“Data Migration Process” section on page 2-28.](#)
- Identify the servers, the WAFS file servers that will be used as the target WAFS file servers, and the desired feature set (for example, disconnected mode and home directories).

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

Site and Network Planning

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

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

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

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

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



Note

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

**Note**

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

This section contains the following topics:

- [Windows Network Integration, page 2-5](#)
- [UNIX Network Integration, page 2-6](#)
- [WAFS-Related Ports in a WAAS Environment, page 2-7](#)
- [Firewalls and Directed Mode, page 2-8](#)

Windows Network Integration

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

- [Data Center WAE Integration, page 2-5](#)
- [Branch WAE Integration, page 2-6](#)

**Note**

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

Data Center WAE Integration

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

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

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

**Note**

User permissions are determined by the existing security infrastructure.

Branch WAE Integration

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

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

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

- To enable all branch WAEs in the specified domain to appear in the Network Neighborhood of users within the same domain, ensure that a Domain Master Browser or local Master Browser is active.
- If DHCP is not used, you must manually add the name and IP address of the branch WAE to the DNS server.
- In Active Directory (AD) environments where nontransparent integration of WAFS is being used, add the Cisco WAFS-cached file server names manually to the AD Computer Catalog. Adding these names (including the default prefix and suffix, if any) enables future integration with AD services such as DFS. If DFS is used, note the AD Site name for the current Edge WAE location and update it in the CIFS section of the Edge WAE configuration. This caveat is not applicable if transparent integration of WAFS is being used.

UNIX Network Integration

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

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

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

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



Note

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

WAFS-Related Ports in a WAAS Environment

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

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

Port 4050

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

Ports 139 and 445

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

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

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

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

- Enable WCCP Version 2 on your routers and branch WAE, as described in the *Cisco Wide Area Application Services Quick Configuration Guide*. Alternatively, you can use inline mode on a branch WAE with a Cisco WAE Inline Network Adapter installed.
- Enable port 445 and disable port 139 on your branch WAEs through the WAAS Central Manager GUI. Click the **Edit** icon next to the Edge WAE in the Manage Devices window and choose **Configure > Acceleration > Legacy Services > File Services > Edge Configuration** from the navigation pane.

Ports 88 and 464

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

Port 50139

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

Firewalls and Directed Mode

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

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

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

For more information about configuring directed mode, see the [“Configuring Directed Mode” section on page 5-17.](#)

About Autoregistration and WAEs

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

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

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



Note

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

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

The encapsulated vendor-specific options field should be encoded as a sequence of code/length/value fields of syntax identical to that of the DHCP options field with the following exceptions:

- a. There should not be a “magic cookie” field in the encapsulated vendor-specific extensions field.
- b. Codes other than 0 or 255 may be redefined by the vendor within the encapsulated vendor-specific extensions field but should conform to the tag-length-value syntax defined in section 2.
- c. Code 255 (END), if present, signifies the end of the encapsulated vendor extensions, not the end of the vendor extensions field. If no code 255 is present, then the end of the enclosing vendor-specific information field is taken as the end of the encapsulated vendor-specific extensions field.

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

**Note**

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

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

- Subnet-mask (option 1)
- Routers (option 3)
- Domain-name (option 15)
- Domain-name-servers (option 6)
- Host-name (option 12)

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

Autoregistration is enabled by default on the first interface of the device. For the FE-511, WAE-511, WAE-512, WAE-611, WAE-612, WAE-674, WAE-7326, WAE-7341, and WAE-7371 models, the first interface is GigabitEthernet 1/0. On an NME-WAE device, autoregistration is enabled on the configured interface.

**Note**

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

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

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

**Note**

Autoregistration is automatically disabled if a static IP address is configured or if interface-level DHCP is configured on the same interface as autoregistration. (See the [“Selecting Static IP Addresses or Using Interface-Level DHCP”](#) section on page 2-10.)

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

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

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

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

Selecting Static IP Addresses or Using Interface-Level DHCP

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

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

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

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

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

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

Identifying and Resolving Interoperability Issues

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

- [Interoperability and Support, page 2-11](#)
- [WAAS and Cisco IOS Interoperability, page 2-12](#)
- [WAAS Compatibility with other Cisco Appliances and Software, page 2-16](#)

Interoperability and Support

[Table 2-1](#) lists the hardware, client, and web browser support for the WAAS software.

Table 2-1 Hardware, Client, Web Browser Support

Hardware support	WAE-511, WAE-512, WAE-611, WAE-612, WAE-674, WAE-7326, WAE-7341, WAE-7371, WAVE-274, WAVE-474, WAVE-574, or an NME-WAE network module that is installed in specific Cisco routers. You must deploy the WAAS Central Manager on a dedicated device.
Client support	The WAAS software running on a branch WAE interoperates with these CIFS clients: Windows 98/NT 4.0/2000/XP/2003.
Web browser support	The WAAS GUIs require Internet Explorer 5.5 or later to run.

This section contains the following topics:

- [Unicode Support for the WAAS GUI Interfaces](#)
- [Unicode Support Limitations](#)

Unicode Support for the WAAS GUI Interfaces

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

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

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

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

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

Unicode Support Limitations

The following are Unicode support limitations:

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

WAAS and Cisco IOS Interoperability

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

- [WAAS Support of the Cisco IOS QoS Classification Feature, page 2-12](#)
- [WAAS Support of the Cisco IOS NBAR Feature, page 2-13](#)
- [WAAS Support of the Cisco IOS Marking, page 2-14](#)
- [WAAS Support of the Cisco IOS Queuing, page 2-14](#)

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

**Note**

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

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

WAAS Support of the Cisco IOS QoS Classification Feature

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

- Source IP address or hostname—Supported under WAAS because the source IP address is preserved by the WAAS device.
- Source TCP/UDP port (or port range)—Supported under WAAS because the source port is preserved by the WAAS device.
- Destination IP address or hostname—Supported under WAAS because the destination IP is preserved by WAAS. WAAS relies on interception at the data center for redirecting traffic to the peer WAAS device.
- Destination TCP/UDP port (or port range)—Supported under WAAS because the destination IP is preserved by WAAS. WAAS relies on interception at the data center for redirecting traffic to the peer WAAS device.
- DSCP/IP precedence (TOS)—Supported under WAAS because WAAS copies the settings of incoming packets on to the outgoing packets from WAAS back to the router. If the packets are not colored at connection establishment time (for TCP packets), there might be a delay in propagating the settings because WAAS does not poll these settings periodically. The packets are eventually colored properly. When packets are not colored they are left uncolored by the WAAS software.

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

WAAS Support of the Cisco IOS NBAR Feature

Unlike a traditional type of classification that is specified through a policy filter that is listed in the [“WAAS Support of the Cisco IOS QoS Classification Feature”](#) section on page 2-12, Network-Based Application Recognition (NBAR) classification needs to consider payload. The classification keeps track of any interceptor that modifies the payload because this modification might cause NBAR to not be able to classify the packets. However, the WAAS software does support NBAR.

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

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

As long as the flow is not identified, NBAR will continue to search for classification in the packets. Because compressed packets will not be classified, this situation can unnecessarily burden the CPU (doing packet inspection). Because of the potential degrade in performance and the slight possibility of correctness issues, we strongly recommend that you use a subinterface or a separate physical interface to connect the WAE to the router (as described in the [“Using Tertiary Interfaces or Subinterfaces to Connect WAEs to Routers”](#) section on page 2-24). When you use a tertiary interface or subinterface to connect the WAE to the router, both the performance and correctness issues are addressed because each packet is processed only once.

6. For dynamic classifications, NBAR maintains a per-flow state. Once certain flows are classified, NBAR does not continue to perform deep packet inspection anymore. However, for other flows (for example, Citrix), NBAR does look at packets continuously because the classification may change dynamically in a flow. Therefore, in order to support all NBAR classifications, it is not sufficient to only poll the DSCP settings of packets incoming to WAAS once per flow; you need to poll periodically to identify flow changes. However, the WAAS system expects packets to appear in the sequence of packets belonging to class C1, followed by a sequence of C2, and so forth, so that a polling method is sufficient to track such dynamic changes.



Note This dynamic classification support requires support for marking DSCP/ToS settings, as specified in the [“WAAS Support of the Cisco IOS QoS Classification Feature”](#) section on page 2-12, as well as the tracking of dynamic changes through polling.

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

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

WAAS Support of the Cisco IOS Marking

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

WAAS Support of the Cisco IOS Queuing

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

WAAS Support of the Cisco IOS Congestion Avoidance

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

WAAS Support of the Cisco IOS Traffic Policing and Rate Limiting

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

WAAS Support of the Cisco IOS Signaling

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

WAAS Support of the Cisco IOS Link-Efficiency Operations

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

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

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

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

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

WAAS and Management Instrumentation

For management instrumentation in the WAAS software, note the following:

- NetFlow is supported. However, depending on where statistics are being obtained, you may see compressed values (statistics on optimized traffic) rather than uncompressed values (statistics for unoptimized traffic).
- You may see statistics on optimized and unoptimized traffic.
- IP Service Level Agreements (SLAs) are supported.
- Full support of policies based on Layer 3 and Layer 4 is provided. Policies based on Layer 7 are partially supported because the first few messages are unoptimized.
- Intrusion Detection System (IDS) is partially supported. The first few messages are unoptimized to allow IDS to detect the intrusive strings.
- Cisco IOS security is partially supported with the exception of features that rely on Layer 5 and above visibility.
- IPsec and SSL VPN is supported.
- Access control lists (ACLs) are supported. IP ACLs on the router take precedence over IP ACLs that are defined on the WAE. For more information, see the [“Access Lists on Routers and WAEs”](#) section on page 2-24.
- VPN is supported if the VPN is deployed after WCCP interception occurs.



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

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

WAAS and MPLS

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

WAAS Compatibility with other Cisco Appliances and Software

If a firewall is placed between the clients and the WAE on one side, and the router on the other side of the firewall, default WCCP redirection does not work. However, if there is a router inside the firewall and another router outside the firewall, default WCCP-based redirection does work and WAAS is supported. You can also enable directed mode to avoid firewall traversal issues. For more information, see the [“Configuring Directed Mode”](#) section on page 5-17.

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

- If you have already deployed ACNS in your network, you can also deploy WAAS.

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

WAAS Devices and Device Mode

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

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

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

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

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

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

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

```
waas-cm# configure
waas-cm(config)#
waas-cm(config)# primary-interface gigabitEthernet 1/0
waas-cm(config)# device mode central-manager
waas-cm(config)# exit
waas-cm# copy run start
waas-cm# reload
Proceed with reload?[confirm] y
Shutting down all services, will Reload requested by CLI@ttyS0.
Restarting system.
```

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



Note

You cannot configure a WAE network module (in the NME-WAE family of devices) to operate in WAAS Central Manager mode.

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

Calculating the Number of WAAS Devices Needed

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

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

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

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

- **Number of WAAS Central Managers**—All networks must have at least one WAAS Central Manager. For larger networks, you should consider deploying two WAAS Central Managers for active and standby back up, high availability, and failover. A WAAS Central Manager is deployed on a dedicated appliance.
- **Number of WAEs**—A minimum of two WAEs are required for flow optimization; one WAE is required on either side of a network link (for example, one in the branch office and one in the data center). A single site can have more than one WAE for redundancy purposes.
- **Number of branch WAEs**—At least one branch WAE is required in each remote office. Larger offices usually have multiple departments whose users work with different servers in the central office. In this situation, you can manage your system easier by following the organizational structure with a branch WAE for each department. In certain situations, multiple branch WAEs can be clustered and configured using DFS or WCCP Version 2, providing failover capabilities. WCCP Version 2 is the recommended method for larger user populations.
- **Number of data center WAEs**—Each organization must have at least one data center WAE.

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

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



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

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

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

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

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

Supported Methods of Traffic Redirection

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

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

This section contains the following topics:

- [Advantages and Disadvantages of Using Inline Interception, page 2-18](#)
- [Advantages and Disadvantages of Using WCCP-Based Routing, page 2-19](#)
- [Advantages and Disadvantages of Using PBR, page 2-20](#)
- [Configuring WCCP or PBR Routing for WAAS Traffic, page 2-20](#)

For detailed information about how to configure traffic interception for your WAAS network, see [Chapter 4, “Configuring Traffic Interception.”](#)

Advantages and Disadvantages of Using Inline Interception

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

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

The Cisco WAE Inline Network Adapter contains one or two pairs of LAN/WAN Ethernet ports. If the adapter has two pairs of ports, it can connect to two routers if the network topology requires it.

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

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

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

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

For more information, see the [“Using Inline Mode to Transparently Intercept TCP Traffic”](#) section on page 4-41.

Advantages and Disadvantages of Using WCCP-Based Routing

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

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

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



Note

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

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

The WCCP Version 2 protocol also has a set of attractive features built-in, for example, automatic failover and load balancing between multiple devices. The WCCP-enabled router monitors the liveness of each WAE that is attached to it through the WCCP keepalive messages. If a WAE goes down, the router stops redirecting packets to the WAE. When you use WCCP Version 2, the branch WAE is not made a single point of failure for the WAAS services. The router can also load balance the traffic among a number of branch WAEs.

You can use CLI commands to configure basic WCCP on both the routers and the WAEs, or you can use CLI commands to configure the router for WCCP and use the WAAS Central Manager GUI to configure basic WCCP on the WAEs. In the configuration example provided in the *Cisco Wide Area Application Services Quick Configuration Guide*, the CLI is used to configure basic WCCP on the WAEs.

We recommend that you use the WAAS CLI to complete the initial basic configuration of WCCP on your first branch WAE and data center WAE, as described in *Cisco Wide Area Application Services Quick Configuration Guide*. After you have verified that WCCP transparent redirection is working properly, you can use the WAAS Central Manager GUI to centrally modify this basic WCCP configuration or configure additional WCCP settings (for example, load balancing) for a WAE (or group of WAEs). For more information, see the [“Managing WCCP Configurations for WAEs”](#) section on page 4-12. After you have configured basic WCCP on the router, you can configure advanced WCCP features on the router, as described in the [“Configuring Advanced WCCP Features on a WCCP-Enabled Router”](#) section on page 4-7.

Advantages and Disadvantages of Using PBR

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

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

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

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

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

- PBR does not support load balancing between equal cost routes. Consequently, PBR does not provide scalability for the deployment location.
- PBR is more difficult to configure than WCCP Version 2. For an example of how to configure PBR for WAAS traffic, see the [“Using Policy-Based Routing to Transparently Redirect All TCP Traffic to WAEs”](#) section on page 4-32.

Configuring WCCP or PBR Routing for WAAS Traffic

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

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



Note

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

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

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



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

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

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

This traffic redirection does not use tunneling; the full original quadruple (source IP address, source port number, destination IP address, and destination port number) of the TCP traffic is preserved end to end. The original payload of the TCP traffic is not preserved end to end because the primary function of WAAS is to accelerate WAN traffic by reducing the data that is transferred across the WAN. This change

in payload can potentially impact features on the router (which is performing the WCCP or PBR redirection) that needs to see the actual payload to perform its operation (for example, NBAR). For more information on this topic, see the [“WAAS and Cisco IOS Interoperability” section on page 2-12](#).

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

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

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

**Note**

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

The TCP promiscuous mode service intercepts all TCP traffic that is destined for any TCP port and transparently redirects it to the WAE. The WCCP-enabled router uses service IDs 61 and 62 to access this service.

By default, the IP Protocol 6 is specified for the TCP promiscuous mode service. Consequently, the routers that have been configured to the TCP promiscuous mode service will intercept and redirect all TCP traffic destined for any TCP port to the local WAE. Because the TCP promiscuous mode service is configured on the WAE, the WAE will accept all of the TCP traffic that is transparently redirected to it by specified WCCP routers (for example, Edge-WAE1 will accept all TCP traffic that Edge-Router1 redirects to it). In the branch office, you can intercept packets at the edge LAN and WAN interfaces on the edge routers and redirect the TCP traffic to the local WAE (the branch WAE). In the data center, you can intercept packets at the core LAN and WAN interfaces on the core routers and redirect the TCP traffic to the local WAE (the data center WAE). For more information, see the [“Configuring WAEs as Promiscuous TCP Devices in a WAAS Network” section on page 2-23](#).

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

**Note**

CEF is not required but is recommended for improved performance. WCCP is optimized to make use of IP CEF if CEF is enabled on the router.

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

**Caution**

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

This section contains the following topics:

- [Configuring WAEs as Promiscuous TCP Devices in a WAAS Network](#), page 2-23
- [Using Tertiary Interfaces or Subinterfaces to Connect WAEs to Routers](#), page 2-24

Configuring WAEs as Promiscuous TCP Devices in a WAAS Network

In order for the WAE to function as a promiscuous TCP device for the TCP traffic that is transparently redirected to it by the specified WCCP Version 2 routers, the WAE uses WCCP Version 2 services 61 and 62. The WCCP services 61 and 62 are represented by the canonical name `tcp-promiscuous` on the WAE, as shown in the following sample output of the WAAS CLI on a branch WAE:

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

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

Although you can also use the WAAS Central Manager GUI to configure the TCP promiscuous mode service on an individual WAE, we recommend that you use the WAAS CLI to complete the initial basic configuration of the WAE and then use the WAAS Central Manager GUI to make any subsequent configuration changes. By using the WAAS Central Manager GUI to make subsequent configuration changes, you can also apply those changes to groups of WAEs (device groups). For instructions on how to perform a basic WCCP configuration for a WAAS network, see the *Cisco Wide Area Application Services Quick Configuration Guide*. For instructions about how to use the WAAS Central Manager GUI to modify the basic WCCP configuration for a WAE or group of WAEs, see the [“Managing WCCP Configurations for WAEs”](#) section on page 4-12.

Using Tertiary Interfaces or Subinterfaces to Connect WAEs to Routers

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



Note

The WCCP GRE return and generic GRE egress methods allow you to place WAEs on the same VLAN or subnet as clients and servers. For information on configuring these egress methods, see the [“Configuring Egress Methods for Intercepted Connections”](#) section on page 4-29.

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

To avoid an infinite routing loop between the router and its local WAE, connect the WAE to the router through a tertiary interface (a separate physical interface) or a subinterface (a different virtual subinterface) from the router's LAN port. By using a tertiary interface or a subinterface to connect a WAE to the router that is performing the PBR or WCCP redirection, the WAE has its own separate processing path that has no Cisco IOS features enabled on it. In addition, this approach simplifies the process of integrating WAEs into an existing network. Because the WAEs are being connected to the routers through a tertiary interface or subinterface that has no Cisco IOS features enabled on it, the Cisco IOS features that are already enabled on your existing Cisco-enabled network elements (for example, Edge-Router1 or Core-Router1) will generally not be affected when you connect WAEs to these routers. For more information about WAAS and Cisco IOS interoperability, see the [“WAAS and Cisco IOS Interoperability” section on page 2-12](#).

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

Access Lists on Routers and WAEs

You can optionally configure the router to redirect traffic from your WAE based on access lists that you define on the router. These access lists are also referred to as redirect lists. For information about how to configure access lists on routers that will be configured to transparently redirect traffic to a WAE, see the [“Configuring IP Access Lists on a Router” section on page 4-10](#).

**Note**

IP access lists on routers have the highest priority followed by IP ACLs that are defined on the WAEs.

This section contains the following topics:

- [IP ACLs on WAEs, page 2-25](#)
- [Static Bypass Lists on WAEs, page 2-25](#)

IP ACLs on WAEs

In a centrally managed WAAS network environment, administrators need to be able to prevent unauthorized access to various devices and services. The WAAS software supports standard and extended IP access control lists (ACLs) that allow you to restrict access to or through a WAAS device. You can use IP ACLs to reduce the infiltration of hackers, worms, and viruses that can harm the corporate network. For example, you can configure IP ACLs on a WAAS device for inbound WCCP encapsulated traffic.

The WAAS software also provides controls that allow various services to be associated with a particular interface. For example, you can use IP ACLs to define a public interface on the WAE for file serving and a private interface for management services, such as SNMP. For more information, see [Chapter 8, “Creating and Managing IP Access Control Lists for WAAS Devices.”](#)

**Note**

IP ACLs that are defined on a WAE always take precedence over any WAAS application definitions that have been defined on the WAE.

Static Bypass Lists on WAEs

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

You can use static bypass occasionally when you want to prevent WAAS from caching a connection from a specific client to a specific file server (or from a specific client to all file servers). For information about how to centrally configure static bypass lists for a WAE or a group of WAEs, see the [“Configuring Static Bypass Lists for WAEs”](#) section on page 4-28.

**Note**

We recommend that you use access lists on the WCCP-enabled router, rather than using static bypass lists on the WAEs, because access lists are more efficient. For information about how to configure access lists on a router, see the [“Configuring IP Access Lists on a Router”](#) section on page 4-10.

WAAS Login Authentication and Authorization

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

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

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

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

For more information, see the [“Configuring AAA Accounting for WAAS Devices”](#) section on page 6-30.

WAAS Administrator Accounts

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

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



Note

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

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

After you have initially configured your WAAS devices, we strongly recommend that you immediately change the password for the predefined superuser account (the predefined username is `admin`, the password is `default`, and the privilege level is superuser, privilege level 15) on each WAAS device. For instructions on how to use the WAAS Central Manager GUI to change the password, see the [“Changing the Password for Your Own Account”](#) section on page 7-7.

Logically Grouping Your WAEs

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

With the WAAS Central Manager GUI, you can easily organize your branch and data center WAEs into the following types of device groups:

- Standard Device Group—A collection of WAEs that share common qualities and capabilities. Setting up groups based on their authentication settings is an example of a device group. There are two types of device groups:
 - Configuration Groups
 - Wide Area File Services (WAFS) Core Clusters (only used for WAFS legacy mode)

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

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

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



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

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

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



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

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

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

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

Data Migration Process

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

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

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

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

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

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

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