CHAPTER 4

Configuring Traffic Interception

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

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

Before you do the procedures in this chapter, you should complete a basic initial installation and configuration of your WAAS network as described in the *Cisco Wide Area Application Services Quick Configuration Guide*. For detailed command syntax information for any of the CLI commands in this chapter, see the *Cisco Wide Area Application Services Command Reference*. For more information about WCCP, see the *Cisco IOS Configuration Fundamentals Configuration Guide* and the *Cisco IOS Configuration Fundamentals Command Reference*.

This chapter contains the following sections:

Request Redirection Methods

In a WAAS network, traffic between clients in the branch offices and the servers in the data center can be redirected to WAEs for optimization, redundancy elimination, and compression. Traffic is intercepted and redirected to WAEs based on policies that have been configured on the routers. The network elements that transparently redirect requests to a local WAE can be a router using WCCP Version 2 or

PBR to transparently redirect traffic to the local WAE or a Layer 4 to Layer 7 switch (for example, the Catalyst 6500 series Content Switching Module [CSM] or Application Control Engine [ACE]). Alternately, you can intercept traffic directly by using the inline mode with a WAE that has a Cisco WAE Inline Network Adapter or Interface Module.

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

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

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

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

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

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

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

  VPATH intercepts traffic from the VM server and redirects it to a vWAAS device for WAN optimization.

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

Table 4-1 Supported Methods of Transparent Traffic Interception

Method	Comment
WCCP Version 2	Used for transparent interception of application traffic and Common Internet File System (CIFS) traffic. Used in branch offices and data centers to transparently redirect traffic to the local WAE. The traffic is transparently intercepted and redirected to the local WAE by a WCCP-enabled router or a Layer 3 switch. You must configure WCCP on the router and WAE in the branch office and the router and WAE in the data center. For more information, see the "Using WCCP to Transparently Redirect TCP Traffic to WAEs" section on page 4-3.
PBR	In branch offices, used for wide area application optimization. The branch office router is configured to use PBR to transparently intercept and route both client and server traffic to the WAE that resides in the same branch office. In data centers, used for data center application optimization. The data center router or L3 switch may be configured to use PBR to transparently intercept and route client and server traffic to WAE(s) within the data center. PBR, however, does not support load balancing across multiple WAEs (such as WCCP does). Neither does it support load balancing when you are using a hardware load balancer, such as the Cisco CSM or ACE. See the "Using Policy-Based Routing to Transparently Redirect All TCP Traffic to WAEs" section on page 4-32.
Inline	Used for transparent interception of application traffic and CIFS traffic. See the "Using Inline Mode to Transparently Intercept TCP Traffic" section on page 4-41.
VPATH	Used for VPATH interception on vWAAS devices. See the "Configuring VPATH Interception on a vWAAS Device" section on page 4-51.
ACE or CSM	Cisco Application Control Engine (ACE) or Catalyst 6500 series Content Switching Module (CSM) installed in the data center for data center application optimization. The ACE or CSM allows for both traffic interception and load balancing across multiple WAE(s) within the data center.

If a WAE device is behind a firewall that prevents traffic optimization, you can use the directed mode of communicating between peer WAEs over the WAN. For details, see the "Configuring Directed Mode" section on page 5-23.

Using WCCP to Transparently Redirect TCP Traffic to WAEs

The WAAS software uses the WCCP standard, Version 2 for redirection. The main features of WCCP Version 2 include support for the following:

- Up to 32 WAEs per WCCP service
- Multiple routers
- Multicasting of protocol messages between the WAE and the WCCP-enabled router
- Authentication of protocol packets
- Redirection of non-HTTP traffic
- Packet return (including GRE, allowing a WAE to reject a redirected packet and to return it to the router to be forwarded)
- Layer 2-caching (through router versus GRE) and masking (for improved load balancing)

- Multiple forwarding methods
- Packet distribution method negotiation within a service group
- Command and status interaction between the WAE and a service group

Note WCCP works only with IPv4 networks.

WAAS software supports the WCCP TCP promiscuous mode service (services 61 and 62 by default, though these service IDs are configurable). This WCCP service requires that WCCP Version 2 is running on the router and the WAE.

The TCP promiscuous mode service is a WCCP service that intercepts all TCP traffic and redirects it to the local WAE.

The WAAS software also supports service passwords, WAE failover, flow protection, interception ACLs, and static bypass.

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

Note Many legacy Cisco routers, including the 2500, 2600, and 3600 routers, have far less processing power and memory than newer routing platforms such as the Integrated Services Router (ISR) models 2800 and 3800. As such, the use of WCCPv2 or PBR may cause a high level of CPU utilization on the router and cause erratic behavior. WAAS can be configured to work with these routers, but not to the same levels of performance or scalability as can be found with newer routing platforms. The Cisco ISR is the routing platform of choice for the branch office.

If you are experiencing erratic behavior, such as the WAE being ejected from the service group, enable fair-queuing, weighted fair-queuing, or rate-limiting on all physical interfaces on the router that connect to users, servers, WAEs, and the WAN. Fair-queuing cannot be configured on subinterfaces, and should be configured on both ingress and egress physical interfaces. If another form of queuing is already configured on the LAN or WAN interfaces other than fair-queuing that provides similar fairness, it should be sufficient.

Additionally, limit the amount of bandwidth that can be received on the LAN-side interface of the router, to help the router keep its interface queues less congested and provide better performance and lower CPU utilization. Set the maximum interface bandwidth on the router to no more than 10 times the WAN bandwidth capacity. For instance, if the WAN link is a T1, the LAN interface and WAE LAN interface bandwidth should be throttled to 10 * T1 = 10 * 1.544 Mbps, or approximately 15 Mbps. See the Cisco IOS documentation for more information.

This section contains the following topics:

Guidelines for Configuring WCCP

When you configure transparent redirection on a WAE using WCCP Version 2, follow these guidelines:

- Intercept and redirect packets on the inbound interface whenever possible.
- Use WCCP GRE or generic GRE as the egress method if you want to place WAEs on the same VLAN or subnet as clients and servers. This topology is not allowed when using the IP forwarding egress method.
- Branch WAEs must not have their packets encrypted or compressed and should be part of the "inside" Network Address Translation (NAT) firewall if one is present.
- Use Layer 2 redirection as the packet forwarding method if you are using Catalyst 6500 series switches or Cisco 7600 series routers. Use Layer 3 GRE packet redirection if you are using any other Cisco series router.
- When you configure WCCP for use with the Hot Standby Router Protocol (HSRP), you must configure the WAE with the HSRP or the Virtual Router Redundancy Protocol (VRRP) virtual router address as its default gateway, and the WAE WCCP router-list with the primary address of the routers in the HSRP group.
- CEF is required for WCCP and must be enabled on the router.
- Place branch WAEs on the client side of the network to minimize client-side packets through the router.
- Use WCCP passwords to avoid denial-of-service attacks. For more information, see the "Setting a Service Group Password on a Router" section on page 4-11.
- Use WCCP redirect lists for new implementations to limit client or server populations. For more information, see the "Configuring IP Access Lists on a Router" section on page 4-10.
- You must configure the WAE to accept redirected packets from one or more WCCP-enabled routers.
- You can quickly view a list of WCCP settings and services that you can configure on a WAE, from the WAAS CLI or the WAAS Central Manager GUI or the WAAS CLI. From the WAAS CLI, enter the **wccp** EXEC command followed by a question mark (?). The following sample output is from a WAE with WCCP Version 2 enabled:

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

- To configure basic WCCP, you must enable the WCCP service on at least one router in your network and on the WAE that you want the traffic redirected to. It is not necessary to configure all of the available WCCP features or services to get your WAE up and running. For an example of how to complete a basic WCCP configuration on routers and WAEs in a branch office and data center, see the *Cisco Wide Area Application Services Quick Configuration Guide.*
- You must configure the routers and WAEs to use WCCP Version 2 instead of WCCP Version 1 because WCCP Version 1 only supports web traffic (port 80).
- After enabling WCCP on the router, you must configure the TCP promiscuous mode service on the router and the WAE, as described in the *Cisco Wide Area Application Services Quick Configuration Guide.* The service IDs are configurable on the WAE and you can choose a pair of numbers different

from the default of 61 and 62, to allow the router to support multiple WCCP farms because the WAEs in different farms can use different service IDs. The router configuration must use WCCP service IDs that match those configured on the WAEs in each farm that it is supporting.

- In order for the WAE to function in TCP promiscuous mode, the WAE uses WCCP Version 2 services 61 and 62 (the service IDs are configurable). These two WCCP services are represented by the canonical name tcp-promiscuous on the WAE.
- You can use CLI commands to configure basic WCCP on both the routers and the WAEs, or you can use CLI commands to configure the router for WCCP and use the WAAS Central Manager GUI to configure basic WCCP on the WAEs. In the configuration example provided in the *Cisco Wide Area Application Services Quick Configuration Guide*, the CLI is used to configure basic WCCP on the WAEs.

  We recommend that you use the WAAS CLI to complete the initial basic configuration of WCCP on your first branch WAE and data center WAE, as described in the *Cisco Wide Area Application Services Quick Configuration Guide.*

  After you have verified that WCCP transparent redirection is working properly, you can use the WAAS Central Manager GUI to centrally modify this basic WCCP configuration or configure additional WCCP settings (for example, load balancing) for a WAE (or group of WAEs). For more information, see the "Managing WCCP Configurations for WAEs" section on page 4-12.
- After you have configured basic WCCP on the router, you can configure advanced WCCP features on the router, as described in the "Configuring Advanced WCCP Features on a WCCP-Enabled Router" section on page 4-7.
- To ensure consistency, we recommend that you configure WCCP settings on a device group basis instead of on an individual device, and the device group should contain WAEs only from a single WCCP service farm. It is best not to configure WCCP using a device group that contains WAEs that belong to multiple farms, since WCCP settings may need to be different in different farms.
- When you add a new router to an existing WCCP router farm or WCCP service group, the new router will reset existing connections. Until WCCP reestablishes path redirections and assignments, packets are sent directly to the client (as expected).
- The router must support the redirect and return methods configured on the WAE. If the router does not support the configured methods, the WAE will not join the WCCP router farm. If you have a mix of routers in the farm, only those routers that support the configured methods will join the farm.
- The WAE only joins the WCCP farm if the assignment method configured on the WAE is supported by the router. (The strict assignment method is always enforced with version 4.4.1 and later.)
- A WAE joins a WCCP farm only if it is seen by all the configured routers in the farm. If there is a link failure with any one of the routers, the farm reconfigures and the WAE is removed from the farm.
- All WAEs in a WCCP farm must use the same pair of WCCP service IDs (the default is 61 and 62), and these IDs must match all routers that are supporting the farm. A WAE with different WCCP service IDs is not allowed to join the farm and an alarm is raised. Likewise, all WAEs in a farm must use the same value for the failure detection timeout. A WAE raises an alarm if you configure it with a mismatching value.
- When you register a WAE to a Central Manager, the WAE is likely to be automatically assigned to a device group, such as the AllWAASGroup. If the WCCP service ID settings of the WAE do not match the service IDs configured for the device group, the WAE retains its previously configured service ID settings and the device group settings are not applied. The WAE appears in override mode in the Central Manager after one or two datafeed poll cycles. You may have to apply the device group service ID settings manually from the Central Manager after the WAE WCCP service ID settings are synchronized to the Central Manager.

- VPN routing and forwarding (VRF) aware WCCP scalability is as follows:
  - The maximum number of WAEs supported by a single VRF instance is 32.
  - The maximum number of VRF instances supported by the router is router dependent.
  - VRF aware WCCP is supported only on specific versions of IOS. Ensure that the router is running a version of IOS that supports VRF aware WCCP.
  - Each VRF instance has independent assignment, redirection, and return methods.

Guidelines for File Server Access Methods

Some file servers have several network interfaces and can be reached through multiple IP addresses. For these server types, you must add all the available IP addresses to the branch WAE's WCCP accept list. This situation prevents a client from bypassing the branch WAE by using an unregistered IP address. The WAE Device Manager GUI displays all the IP addresses in the GUI.

Some file servers have several NetBIOS names and only one IP address. For these servers, if the client connects using the IP address in the UNC path (that is, \\IP_address\share instead of \\server\share), WAAS selects the first NetBIOS name from the server list in the WAE Device Manager GUI that matches this IP address. WAAS uses that name to perform NetBIOS negotiations between the data center WAE and the file server, and to create resources in the cache. If a file server uses multiple NetBIOS names to represent virtual servers (possibly with different configurations) and has one NetBIOS name that is identified as the primary server name, put that name in the server list before the other names.

Configuring Advanced WCCP Features on a WCCP-Enabled Router

This section describes how to configure the advanced WCCP Version 2 features on a WCCP-enabled router that is transparently redirecting requests to WAEs in your WAAS network and contains the following topics:

Note Before you do the procedures in this section, you should have already configured your router for basic WCCP as described in the *Cisco Wide Area Application Services Quick Configuration Guide.*

Configuring a Router to Support WCCP Service Groups

WCCP Version 2 enables a set of branch WAEs in a WAE group to connect to multiple routers. The WAEs in a group and the WCCP Version 2-enabled routers connected to the WAE group that are running the same WCCP service are known as a *service group*.

Through communication with the branch WAEs, the WCCP Version 2-enabled routers are aware of the available branch WAEs. Routers and branch WAEs become aware of one another and form a service group using WCCP Version 2. See Figure 4-1.

Figure 4-1 Service Groups with WCCP Version 2

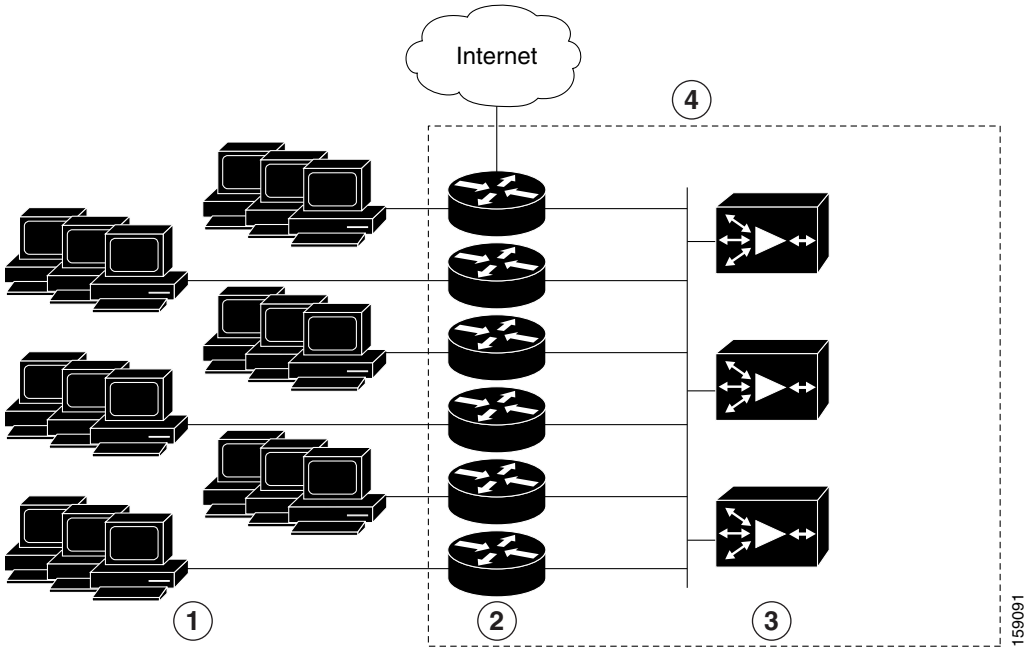

1	Clients requesting file services	**3**	Branch WAEs
2	Cisco routers	**4**	WAE service group

If you have a group of branch WAEs, the WAE that is seen by all the WCCP Version 2-enabled routers and that has the lowest IP address becomes the lead branch WAE.

The following procedure describes how a branch WAE in a service group is designated as the lead:

1. Each branch WAE is configured with a list of WCCP-enabled routers.

   Multiple WCCP-enabled routers can service a group (up to 32 routers can be specified). Any of the available routers in a service group can redirect packets to each of the branch WAEs in the group.

2. Each branch WAE announces its presence to each router on the router list. The routers reply with their view of branch WAEs in the service group.

3. After the view is consistent across all of the branch WAEs in the group, one branch WAE is designated as the lead branch WAE and sets the policy that the WCCP-enabled routers need to deploy in redirecting packets.

The lead branch WAE determines how traffic should be allocated across the branch WAEs in the group. The assignment information is passed to the entire service group from the designated lead branch WAE so that the WCCP-enabled routers of the group can redirect the packets and the branch WAEs in the group can better manage their load.

WCCP uses service groups to define WAAS services for a WCCP Version 2-enabled router and branch WAEs in a group. WCCP also redirects client requests to these groups in real time.

All ports receiving redirected traffic that are configured as members of the same WCCP service group share the following characteristics:

- They have the same hash or mask parameters, as configured with the WAAS Central Manager GUI (the "Viewing or Modifying the WCCP Settings on WAEs" section on page 4-18) or the WAAS CLI (the **wccp service-number mask** global configuration command).
- The WCCP Version 2 service on individual ports cannot be stopped or started individually (a WCCP Version 2 restriction).

To direct a WCCP Version 2-enabled router to enable or disable support for a WCCP service group, use the **ip wccp** global configuration command. To remove the ability of a router to control support for a WCCP service group, use the **no** form of this command.

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

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

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

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

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

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

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

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

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

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

Finally, you need to configure each router for WCCP interception on the inbound direction of the appropriate interfaces, using commands similar to the following:

```
Router(config)# interface fa1/0.40
Router(config-subif)# ip wccp 61 redirect in
Router(config-subif)# exit
Router(config)# interface serial0
Router(config-subif)# ip wccp 62 redirect in
Router(config-subif)# exit
```

When a new WAE is brought online, it joins the WCCP service group. With a new WAE in the service group, the hash tables responsible for distributing the load are changed, and traffic that previously went to WAE1 may now go to WAE2. Flow protection must be enabled in order for WAE2 to forward packets of already connected clients to WAE1. The end result is that all requests that belong to a single session are processed by the same WAE. Should the administrator disable flow protection, adding a WAE to the service group might disconnect some of the existing clients.

When an WAE is removed from the service group, its clients are disconnected (if they reconnect, they will reach another WAE, if one is available, or the origin file server).

WAAS supports WAE failover by reconnecting clients with other branch WAEs if a branch WAE crashed. In the event of a crash, the branch WAE stops issuing WCCP keepalives (constant high CPU load may also result in loss of keepalives and can also be considered a failover case). The router detects the lack of keepalives and removes the branch WAE from the service group. The designated branch WAE updates the WCCP configuration hash table to reflect the loss of the branch WAE and divides its buckets among the remaining branch WAEs. A new designated lead branch WAE is elected if the crashed one was the lead branch WAE. The client is disconnected, but subsequent connections are processed by another branch WAE.

Once a TCP flow has been intercepted and received by a branch WAE, the failure behavior is identical to that exhibited during nontransparent mode. For example, data center WAE and file server failure scenarios are not handled any differently as a result of using WCCP interception.

Note When you add a new router to an existing WCCP router farm or WCCP service group, the new router will reset existing connections. Until WCCP reestablishes path redirections and assignments, packets are sent directly to the client (as expected).

Configuring IP Access Lists on a Router

You can optionally configure the router to redirect traffic from your WAE based on access control lists (ACLs) that you define on the router. These access lists are also referred to as redirect lists.

Note We recommend that you use redirect lists on the WCCP-enabled router where possible, because that is the most efficient method to control traffic interception. However, you can also configure static bypass lists or interception ACLs on the WAEs, and of these two, we recommend using interception ACLs because they are more flexible and give better statistics about passed-through connections. For information about how to configure an interception ACL for a WAE, see the "Configuring an Interception Access Control List" section on page 4-27. For information about how to configure a static bypass list, see the "Configuring Static Bypass Lists for WAEs" section on page 4-26. You can also configure interface ACLs on WAEs to control management access to the WAE, as described in Chapter 8, "Creating and Managing IP Access Control Lists for WAAS Devices."

Redirect lists that are configured on the routers have the highest priority, followed by static bypass lists or interception ACLs on WAEs. Interception ACLs that are configured on WAEs take precedence over any application definition policies that have been defined on the WAE.

A WCCP Version 2-enabled router can be configured with access lists to permit or deny redirection of TCP traffic to a WAE. The following example shows that traffic conforming to the following criteria are not redirected by the router to the WAE:

- Originating from the host 10.1.1.1 destined for any other host
- Originating from any host destined for the host 10.255.1.1

```
Router(config)# ip wccp 61 redirect-list 120
Router(config)# ip wccp 62 redirect-list 120
Router(config)# access-list 120 deny ip host 10.1.1.1 any
Router(config)# access-list 120 deny ip any host 10.1.1.1
Router(config)# access-list 120 deny ip any host 10.255.1.1
Router(config)# access-list 120 deny ip host 10.255.1.1 any
```

```
Router(config)# access-list 120 permit ip any
```

Traffic not explicitly permitted is implicitly denied redirection. The **access-list 120 permit ip any** command explicitly permits all traffic (from any source on the way to any destination) to be redirected to the WAE. Because criteria matching occurs in the order in which the commands are entered, the global **permit** command is the last command entered.

To limit the redirection of packets to those packets matching an access list, use the **ip wccp redirect-list** global configuration command. Use this command to specify which packets should be redirected to the WAE.

When WCCP is enabled but the **ip wccp redirect-list** command is not used, all packets matching the criteria of a WCCP service are redirected to the WAE. When you specify the **ip wccp redirect-list** command, only packets that match the access list are redirected.

The **ip wccp** global configuration command and the **ip wccp redirect** interface configuration command are the only commands required to start redirecting requests to the WAE using WCCP. To instruct an interface on the WCCP-enabled router to check for appropriate outgoing packets and redirect them to a WAE, use the **ip wccp redirect** interface configuration command. If the **ip wccp** command is enabled but the **ip wccp redirect** command is disabled, the WCCP-enabled router is aware of the WAE but does not use it.

To specify the access list by name or number, use the **ip wccp group-list** global configuration command, which defines criteria for group membership. In the following example, the **access-list 1 permit 10.10.10.1** command is used to define the IP address of the WAE that is allowed to join the WCCP service group:

```
Router(config)# ip wccp 61 group-list 1
Router(config)# ip wccp 62 group-list 1
Router(config)# access-list 1 permit 10.10.10.1
```

Tip If you have a WCCP service farm with multiple WAEs, the load balancing assignment may cause packets that are sent to the WAE devices themselves (such as management traffic) to be redirected to a different WAE in the farm, negatively impacting performance. To avoid this situation, we recommend that you configure a WCCP redirect list that excludes traffic that is sent to the WAE IP addresses from being redirected.

For more information on access lists, see the Cisco IOS IP addressing and services software documentation.

Setting a Service Group Password on a Router

For security purposes, you can set a service password for your WCCP Version 2-enabled router and the WAEs that access it. Only devices configured with the correct password are allowed to participate in the WCCP service group.

From the global configuration mode of your WCCP-enabled router, enter the following commands to specify the service group password for the TCP promiscuous mode service on the router (the service IDs must match the service IDs configured on the WAE):

```
Router(config)# ip wccp 61 password [0-7] password
Router(config)# ip wccp 62 password [0-7] password
```

The required *password* argument is the string that directs the WCCP Version 2-enabled router to apply MD5 authentication to messages received from the specified service group. Messages that are not accepted by the authentication are discarded. *0-7* is the optional value that indicates the HMAC MD5 algorithm used to encrypt the password. This value is generated when an encrypted password is created for the WAE. 7 is the recommended value. The optional *password* argument is the optional password name that is combined with the HMAC MD5 value to create security for the connection between the router and the WAE.

For information about how to use the WAAS Central Manager GUI to specify the service group password on a WAE (or device group), see the "Viewing or Modifying the WCCP Settings on WAEs" section on page 4-18.

Configuring a Loopback Interface on the Router

The highest IP address among the router's loopback interfaces is used to identify the router to the WAEs.

The following example configures the loopback interface, exits configuration mode, and saves the running configuration to the startup configuration:

```
Router(config)# interface Loopback0
Router(config-if)# ip address 111.111.111.111 255.255.255.0
Router(config-if)# no shutdown
Router(config-if)# end
Router# copy running-config startup-config
```

Configuring Router QoS for WCCP Control Packets

Beginning with version 4.2.1, WAAS sends WCCP control packets marked with a differentiated services code point (DSCP) value of 192. (Previously, packets were unmarked.) For a router to honor this priority value, you must configure the router's multilayer switching (MLS) quality of service (QoS) port trust state and classify traffic by examining the DSCP value. To configure the router appropriately, use the **mls qos trust dscp** command in interface configuration mode on the interface connected to the WAE.

Managing WCCP Configurations for WAEs

This section contains the following topics:

- Configuring an Interception Access Control List, page 4-27

Note Before you do the procedures in this section, you should have completed an initial configuration of your WAAS network, which includes the basic configuration of WCCP Version 2 and the TCP promiscuous mode service on your routers and WAEs, as described in the *Cisco Wide Area Application Services Quick Configuration Guide*.

Load Balancing and WAEs

Multiple WAEs with WCCP support can be deployed for dynamic load balancing to enable adjustments to the loads being forwarded to the individual WAEs in a service group. IP packets received by a WCCP-enabled router are examined to determine if it is a request that should be directed to a WAE. Packet examination involves matching the request to a defined service criteria. These packets are passed to the processing routine on the router to determine which WAE, if any, should receive the redirected packets.

You can use load balancing to balance the traffic load across multiple WAEs. Load balancing allows the set of hash address buckets assigned to a WAE to be adjusted, shifting the load from an overwhelmed WAE to other WAEs that have available capacity. Two assignment methods are used by this technique: hashing and masking.

The term *assignment method* denotes the method used by WCCP to perform load distribution across WAEs. The two possible load-balancing assignment methods are hashing and masking. If the mask load-balancing method is not specified, then the hash load-balancing method, which is the default method, is used.

WCCP supports redirection based on a hash function. The hash key may be based on the source or destination IP address of the packet. For WAAS, load-balancing hashing is based on a source IP address (default), a destination IP address, or both.

The hash function uses the source IP address to obtain an address bucket to which the packet is assigned. These source address buckets are then mapped to a particular WAE depending on how many WAEs are present and how busy they are. (See Figure 4-2.)

Figure 4-2 Load Balancing Through Hashing of IP Addresses

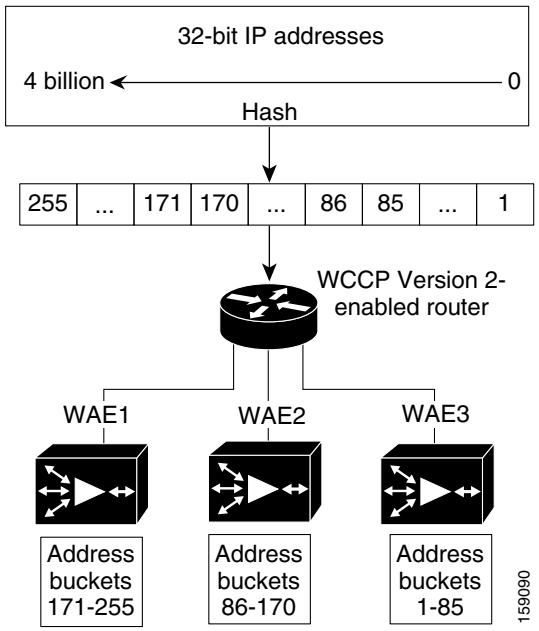

Note Packets that the WAEs do not service are tunneled back to the same router from which they were received. When a router receives a formerly redirected packet, it knows not to redirect it again.

Destination IP address hashing guarantees that a single WAE caches a given file server. This method, which allows a local coherency directive to be safely applied to the file server content (provided that no other collaboration on the content occurs), improves performance and WAN link and disk utilization. This method may distribute load unevenly, however, because of uneven activity on a file server.

Source IP address hashing has better potential for session distribution between the caches on branch WAEs. This method may impact performance and WAN link and disk utilization (see the previous description of factors to be aware of when load balancing is applied). Also, any change in the IP address of a client (which can happen when working in DHCP environments) may cause the client to switch to another branch WAE, which can cause the client to experience reduced performance until the client's working set is retrieved into the new cache.

Hashing that is based on a client IP address does not guarantee any locality of the hash key. For example, clients from the same subnet (which are likely to share and collaborate on the same content) may be assigned two different hash numbers and may be redirected to different branch WAEs, while clients from different subnets may be assigned the same hash number and may be redirected to the same branch WAE. Hashing that is based on a client IP address does guarantee consistency. For example, a client using the same IP address is redirected to the same branch WAE.

In the service farm, a lead WAE is chosen to build the hash table that distributes the load between the available WAEs. The lead WAE distributes the buckets evenly. The source IP address is hashed and the resulting bucket determines the WAE that will handle the packet (flow protection tries to ensure that it is the same WAE throughout the session).

WCCP supports redirection by mask value assignments. This method relies on masking to make redirection decisions. The decisions are made using special hardware support in the WCCP-enabled router. This method can be very efficient because packets are switched by the hardware.

Note The masking method can only be used for load balancing with the Catalyst 3750, Catalyst 4500, and Catalyst 6500 series switches, Cisco 7600 series routers, and Cisco ASR 1000 series routers. And, the masking method can be used with the Cisco 2800, 3800, and 7200 series routers when they are running IOS version 12.4(20)T or later.

You must explicitly specify masking. You can specify two mask values based on the source or destination IP address of the packet. For WAAS, the default mask value is based on the source IP address. You can enable masks by using the default values or specifying a particular mask. The default mask values, specified in hexadecimal notation, are as follows:

- dst-ip-mask= 0x0
- src-ip-mask= 0xF00

You may specify the mask value with a maximum of seven bits. The WAE creates a table of the 2^7 (or 128) combinations, assigns the WAE IP addresses to them, and sends this table to the WCCP-enabled routers. The router uses this table to distribute the traffic among all the WAEs that are in the service group. Each packet that matches the WCCP service parameters is compared to this table and the packets are sent to the matching WAE.

In a service farm where the WAEs have different masks, the first WAE to establish two-way communication with the router(s) determines the farm's mask. All other WAEs cannot join the farm unless they are configured with the same mask.

Masking is typically used at the data center, where you can take advantage of the hardware accelerated WCCP redirection capabilities of switches such as the Catalyst 6500 series switches. At the data center, the load balancing goal should be to have all connections originating from a given client subnet (typically equivalent to a branch) go to one datacenter WAE, to improve DRE compression performance. Also, mask assignment on the Catalyst 6500 series switches uses the ACL TCAM. When combined with WCCP redirect lists, mask assignment can use a large portion of the TCAM. To minimize TCAM usage, use a mask with fewer care bits.

Given these considerations, beginning with WAAS version 4.2.1, the default mask has been changed from src-ip-mask 0x1741 and dst-ip-mask 0x0 (in 4.1x versions) to src-ip-mask 0xF00 and dst-ip-mask 0x0 (in 4.2.1 and later versions). The current source IP mask uses only 4 care bits rather than the 6 care bits used by the old mask.

With a typical data center WCCP interception configuration (ingress interception with service 61 on the WAN, ingress interception with service 62 on the LAN), this mask load balances /24 branch subnets (it extracts the last 4 bits of /24 subnets). Connections from one branch subnet will be pinned to one data center WAE. If your network has a different distribution of IP addresses (for example, /16 subnets), you should configure a mask that extracts bits from the /16 network part of the address, for example, src-ip-mask 0xF0000. Similarly, if some branches generate more traffic than others, you may want to create a mask that also extracts bits from the host part of the address, for example 0xF03.

Packet-Forwarding Methods

A WCCP-enabled router redirects intercepted TCP segments to a WAE using one of the following two packet-forwarding methods:

- Generic routing encapsulation (GRE)—Allows packets to reach the WAE even if there are any number of routers in the path to the WAE.
- Layer 2 redirection—Allows packets to be switched at Layer 2 (MAC layer) and reach the WAE.

Table 4-2 describes the packet-forwarding methods.

Table 4-2 Packet-Forwarding Methods

Packet-Forwarding Method	Load-Balancing Method: Hashing	Load-Balancing Method: Masking
GRE (Layer 3)	Packet redirection is completely handled by the router software.	Packet redirection is handled by the router software. We do not recommend using mask assignment when GRE is being used as the packet-forwarding method.
Layer 2 redirection	First redirected packet is handled by the router software; all subsequent redirected packets are handled by the router hardware.	All packets are handled by the router hardware (currently supported only on the Catalyst 6500 series switches or Cisco 7600 series routers because special hardware is required).

The redirection mode is controlled by the branch WAE. The first branch WAE that joins the WCCP service group decides the forwarding method (GRE or Layer 2 redirection) and the assignment method (hashing or masking). The term *mask assignment* refers to WCCP Layer 2 Policy Feature Card 2 (PFC2) input redirection.

If masking is selected with WCCP output redirection, then the branch WAE falls back to the original hardware acceleration that is used with the Multilayer Switch Feature Card (MSFC) and the Policy Feature Card (PFC).

For example, WCCP filters the packets to determine which redirected packets have been returned from the branch WAE and which ones have not. WCCP does not redirect the ones that have been returned because the branch WAE has determined that the packets should not be processed. WCCP Version 2 returns packets that the branch WAE does not service to the same router from which they were transmitted.

This section contains the following topics:

Reasons for Packet Rejection and Return

A branch WAE rejects packets and initiates packet return for the following reasons:

- The WAE is filtering out certain conditions that make processing packets unproductive, for example, when IP authentication has been turned on.
- You have configured a static bypass list or interception ACL on the branch WAE.

Note The packets are redirected to the source of the connection between the WCCP-enabled router and the branch WAE. Depending on the Cisco IOS software version used, this source could be either the address of the outgoing interface or the router IP address. In the latter case, it is important that the branch WAE has the IP address of the WCCP-enabled router stored in the router list. For more information on router lists, see the "Modifying the Configuration of WCCP Router Lists for WAEs" section on page 4-23.

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

WCCP also allows you to configure multiple routers (router lists) to support a particular WCCP service (for example, CIFS redirection).

Layer 3 GRE as a Packet-Forwarding Method

A WCCP-enabled router redirects intercepted requests to a WAE and can encapsulate the packets using generic routing encapsulation (GRE). This method for forwarding packets allows packets to reach the WAE even if there are routers in the path to the WAE. Packet redirection is handled entirely by the router software.

GRE allows datagrams to be encapsulated into IP packets at the WCCP-enabled router and then redirected to a WAE (the transparent proxy server). At this intermediate destination, the datagrams are decapsulated and then handled by the WAAS software. If the request cannot be handled locally, the origin server may be contacted by the associated WAE to complete the request. In doing so, the trip to the origin server appears to the inner datagrams as one hop. The redirected traffic using GRE usually is referred to as GRE tunnel traffic. With GRE, all redirection is handled by the router software.

With WCCP redirection, a Cisco router does not forward the TCP SYN packet to the destination because the router has WCCP enabled on the destination port of the connection. Instead, the WCCP-enabled router encapsulates the packet using GRE tunneling and sends it to the WAE that has been configured to accept redirected packets from this WCCP-enabled router.

After receiving the redirected packet, the WAE does the following:

1. Strips the GRE layer from the packet.
2. Decides whether it should accept this redirected packet and process the request for content or deny the redirected packet as follows:
   a. If the WAE decides to accept the request, it sends a TCP SYN ACK packet to the client. In this response packet, the WAE uses the IP address of the original destination (origin server) that was specified as the source address so that the WAE can be invisible (transparent) to the client; it pretends to be the destination that the TCP SYN packet from the client was trying to reach.
   b. If the WAE decides not to accept the request, it reencapsulates the TCP SYN packet in GRE, and sends it back to the WCCP-enabled router. The router understands that the WAE is not interested in this connection and forwards the packet to its original destination (that is, the origin server).

Layer 2 Redirection as a Packet-Forwarding Method

Layer 2 redirection is accomplished when a WCCP-enabled router or switch takes advantage of internal switching hardware that either partially or fully implements the WCCP traffic interception and redirection functions at Layer 2. This type of redirection is currently supported only with the Catalyst 6500 series switches and Cisco 7200 and 7600 series routers. With Layer 2 redirection, the first redirected traffic packet is handled by the router software. The rest of the traffic is handled by the router hardware. The branch WAE instructs the router or switch to apply a bit mask to certain packet fields, which in turn provides a mask result or index mapped to the branch WAE in the service group in the form of a mask index address table. The redirection process is accelerated in the switching hardware, making Layer 2 redirection more efficient than Layer 3 GRE.

Note WCCP is licensed only on the WAE and not on the redirecting router. WCCP does not interfere with normal router or switch operations.

WCCP Flow Redirection on WAEs

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

- WAAS network expansion—When branch WAEs are added to the service group, the newly started branch WAEs receives traffic that was previously processed by a different branch WAE. It forwards the traffic to the relevant branch WAE for continued processing. New connections are processed by the new branch WAE.
- Branch WAE replacement following a failure—When a branch WAE fails, another branch WAE may receive traffic that was previously processed by either that branch WAE or the origin file server. The receiving branch WAE operates according to the previous two use cases.

Without flow protection, established client connections are broken through a TCP RESET in the situations listed earlier. Flow protection applies to all supported WCCP services and cannot be configured on a per-service basis.

Note Designs that require redirected frames to be returned to the originating router are not compatible with the WCCP flow protection feature.

Viewing or Modifying the WCCP Settings on WAEs

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

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

Note Before you do the procedure in this section, you should have already completed a basic WCCP configuration for your WAAS network that includes the configuration of the TCP promiscuous mode service as described in the *Cisco Wide Area Application Services Quick Configuration Guide*.

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

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

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

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

Figure 4-3 WCCP Settings Window

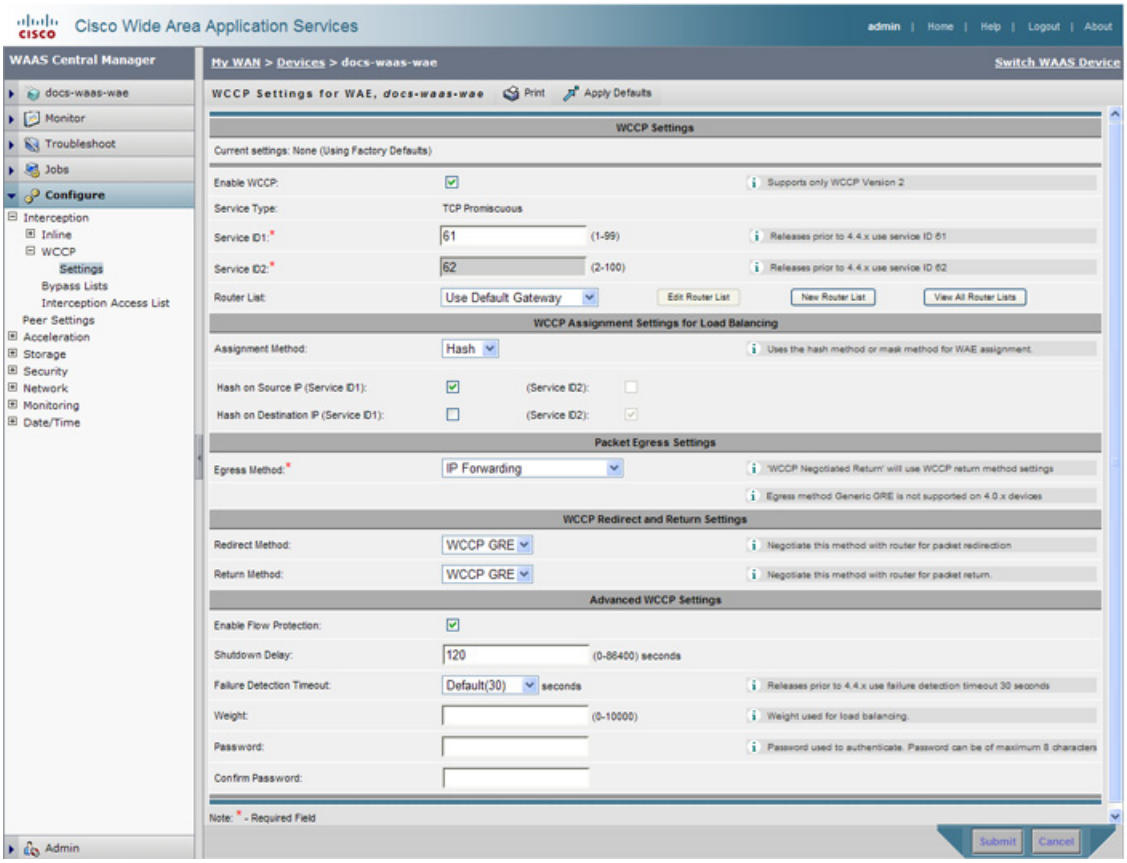

Note For vWAAS devices, only one type of interception can be enabled at a time (WCCP or VPATH). For device groups, WCCP interception is not applicable for vWAAS devices (OE-VWAAS) in the group that have VPATH enabled. Such devices will go to override status.

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

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

By default, WCCP is disabled on a WAE. However, as part of the initial configuration of WCCP in your WAAS network, you should have enabled WCCP Version 2 on your WAEs (the branch WAE and the data center WAE) as well as on the routers in the data center and branch office that will be transparently redirecting requests to these WAEs. For information about how to perform a basic WCCP configuration in your WAAS network, see the *Cisco Wide Area Application Services Quick Configuration Guide*.

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

Note Ensure that the routers used in the WCCP environment are running a version of the Cisco IOS software that also supports the WCCP Version 2.

Step 6 In the Service ID1 field, specify the first service ID of the WCCP service pair. After you submit, the Service ID2 field is filled in with the second service ID of the pair, which is one greater than Service ID1. The default service IDs are 61 and 62, and these IDs are required for any devices that use WAAS versions prior to 4.4.1. On WAAS version 4.4.1 or later, you can change the WCCP service IDs from the default of 61/62 to a different pair of numbers, which allows a router to support multiple WCCP farms because the WAEs in different farms can use different service IDs.

The router service priority varies inversely with the service ID. The service priority of the default service IDs 61/62 is 34. If you specify a lower service ID, the service priority is higher than 34; if you specify a higher service ID, the service priority is lower than 34.

Step 7 Associate a router list with the WCCP TCP promiscuous mode service by choosing the appropriate number of the WCCP router list from the Router List drop-down list. Alternatively, you can choose the default, which is Use Default Gateway. The default router list contains the single IP address of the default gateway of the WAE device.

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

- Modifying the Configuration of WCCP Router Lists for WAEs, page 4-23
- Deleting a WCCP Router List from WAEs, page 4-24
- Defining Additional WCCP Router Lists on WAEs, page 4-24

Step 8 (Optional) From the Assignment Method drop-down list, choose the type of WAE load-balancing assignment method to use (for more information, see the "Load Balancing and WAEs" section on page 4-13):

- Choose **Hash** (the default) to use the hash method. Then follow Steps 9 and 10 to define how the hash works, and skip to Step 12 because the mask settings are not used.
- Choose **Mask** to use the mask method. Skip to Step 11 to define the service mask.

Step 9 (Optional) To define the load-balancing hash for WCCP service ID1 on the source IP address, check the **Hash on Source IP** check box. This check box is shown only if the hash assignment method is used.

Step 10 (Optional) To define the load-balancing hash for WCCP service ID1 on the destination IP address, check the **Hash on Destination IP** check box. This check box is shown only if the hash assignment method is used.

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

Step 11 (Optional) To use a custom mask, uncheck the **Use Default Mask** check box, or leave it checked to use the default mask. (This check box and the mask fields are shown only if the mask assignment method is used.) Define the custom mask as follows:

- In the Source IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, FE000000) used to match the packet source IP address. The range is 00000000–FE000000. The default is F00.
- In the Destination IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, FE000000) used to match the packet destination IP address. The range is 0000000–FE000000. The default is 00000000.

Note If you apply the default mask to a WAE running version 4.1.x or earlier, the mask is different from the default mask (0x1741) set under software version 4.1.x and earlier.

If the WAE detects that its configured mask is not the same as advertised by one or more routers in the farm, it is not allowed to join the farm and a major alarm is raised ("Configured mask mismatch for WCCP"). This alarm can occur when a WAE is trying to join a farm that already has other WAEs and these other WAEs are configured with a different mask. The routers do not allow other WAEs to join the farm unless they advertise the same mask. To correct this alarm, ensure that all WAEs in the farm are configured with the same mask. This alarm is cleared when the WAE's configured mask matches the mask of all the routers in the farm.

Step 12 (Optional) In the Packet Egress Settings area, from the Egress Method drop-down list, choose the type of method to use to return optimized packets to the router or switch: **IP Forwarding** (the default), **WCCP Negotiated Return**, or **Generic GRE**. For more details on choosing the egress method, see the "Configuring Egress Methods for Intercepted Connections" section on page 4-28.

Step 13 (Optional) Modify the current packet redirect and return method settings in the WCCP Redirect and Return Settings area as follows:

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

- Choose **WCCP GRE** (the default) to use Layer 3 GRE packet redirection.
- Choose **WCCP L2** to permit the WAE (or device group) to receive transparently redirected traffic from a WCCP Version 2-enabled switch or router if the WAE has a Layer 2 connection with the device and the device is configured for Layer 2 redirection.

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

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

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

Step 14 (Optional) Modify the current advanced settings in the Advanced WCCP Settings area as follows:

a. Check the **Enable Flow Protection** check box to keep the TCP flow intact and to avoid overwhelming the device (or device groups) when they come up or are reassigned new traffic. For more information, see the "WCCP Flow Redirection on WAEs" section on page 4-18. Flow protection is enabled by default.

b. In the Shutdown Delay field, specify the maximum amount of time (in seconds) that the chosen device (or device group) waits to perform a clean shutdown of WCCP. The default is 120 seconds.

The WAE does not reboot until either all connections have been serviced or the maximum wait time (specified through this Shutdown Delay field) has elapsed for WCCP Version 2.

c. In the Failure Detection Timeout drop-down list, choose the failure detection timeout value (9, 15, or 30 seconds). The default is 30 seconds and is the only value supported on WAAS versions prior to 4.4.1. This failure detection value determines how long it takes the router to detect a WAE failure.

   The failure detection timeout value is negotiated with the router and takes effect only if the router also has the variable timeout capability. If the router has a fixed timeout of 30 seconds and you have configured a failure detection value on the WAE other than the default 30 seconds, the WAE is not able to join the farm and an alarm is raised ("Router unusable" with a reason of "Timer interval mismatch with router").

d. In the Weight field, specify the weight value that is used for load balancing. The weight value ranges from 0 to 10000. If the total of all the weight values of the WAEs in a service group is less than or equal to 100, then the weight value represents a literal percentage of the total load redirected to the device for load-balancing purposes. For example, a WAE with a weight of 10 receives 10 percent of the total load in a service group where the total of all weight values is 50. If a WAE in such a service group fails, the other WAEs still receive the same load percentages as before the failure; they will not receive the load allocated to the failed WAE.

   If the total of all the weight values of the WAEs in a service group is between 101 and 10000, then the weight value is treated as a fraction of the total weight of all the active WAEs in the service group. For example, a WAE with a weight of 200 receives 25 percent of the total load in a service group where the total of all the weight values is 800. If a WAE in such a service group fails, the other WAEs will receive the load previously allocated to the failed WAE. The failover handling is different than if the total weights are less than or equal to 100.

   By default, weights are not assigned and the traffic load is distributed evenly between the WAEs in a service group.

e. In the Password field, specify the password to be used for secure traffic between the WAEs within a cluster and the router for a specified service. Be sure to enable all other WAEs and routers within the cluster with the same password. Passwords must not exceed eight characters in length. Do not use the following characters: space, backwards single quote (`), double quote ("), pipe (|), or question mark (?). Reenter the password in the Confirm Password field.

Note For information about how to use the CLI to specify the service group password on a router, see the "Setting a Service Group Password on a Router" section on page 4-11.

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

For more information about a graceful shut down of WCCP Version 2 on WAEs, see the "Configuring WAEs for a Graceful Shutdown of WCCP" section on page 4-26.

Note When configuring WCCP settings on a device group, if the device group contains devices that run a WAAS version prior to 4.4.1, you are not allowed to configure nondefault values for service IDs or the failure detection timeout. Similarly, devices that run WAAS versions prior to 4.4.1 are not allowed to join a device group that is configured with nondefault values for service IDs or the failure detection timeout. In such cases, we recommend that you upgrade all devices to the same software version or create different device groups for devices with incompatible versions.

Viewing a WCCP Router List Configuration for WAEs

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

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

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

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

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

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

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

Note To modify the configuration of a specific WCCP router list, click the **Edit** icon next to the router list and use the displayed Modifying Router List to modify the chosen router list. For more information about modifying router lists, see the "Modifying the Configuration of WCCP Router Lists for WAEs" section on page 4-23. For information about how to delete a WCCP router list from a WAE (or group of WAEs), see the "Deleting a WCCP Router List from WAEs" section on page 4-24.

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

Modifying the Configuration of WCCP Router Lists for WAEs

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

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

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

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

The WCCP Settings window appears.

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

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

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

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

Deleting a WCCP Router List from WAEs

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

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

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

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

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

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

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

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

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

Defining Additional WCCP Router Lists on WAEs

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

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

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

Note Before you do the procedure in this section, you should have already completed a basic WCCP configuration for your WAAS network that includes the configuration of the TCP promiscuous mode service as described in the *Cisco Wide Area Application Services Quick Configuration Guide*.

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

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

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

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

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

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

Figure 4-4 Creating a New WCCP Router List Window

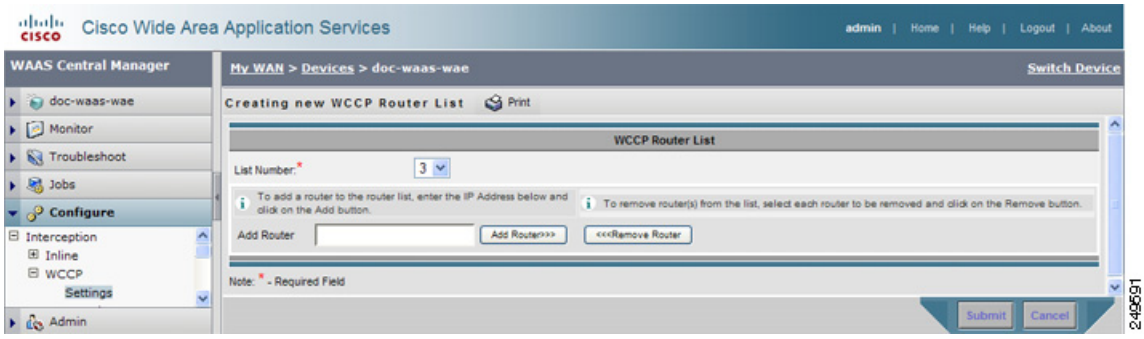

In this example, the number 3 is preselected in the List Number drop-down list because there are already two WCCP router lists defined for the chosen device (or device group). Router list 1 has already been defined for the WCCP router in the data center that will be transparently redirecting traffic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

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

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

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

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

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

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

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

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

Configuring WAEs for a Graceful Shutdown of WCCP

To prevent broken TCP connections, the WAE performs a clean shutdown of WCCP after you disable WCCP Version 2 on a WAE or reload the WAE.

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

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

The WAE does not reboot until one of the following occurs:

- All the connections have been serviced.
- The maximum wait time (specified through the Shutdown Delay field in the WCCP Configuration Settings window or with the **wccp shutdown max-wait** command [by default, 120 seconds]) has elapsed for WCCP Version 2.

During a clean shutdown of WCCP, the WAE continues to service the flows that it is handling, but it starts to bypass new flows. When the number of flows goes down to zero, the WAE takes itself out of the group by having its buckets reassigned to other WAEs by the lead WAE. TCP connections can still be broken if the WAE crashes or is rebooted without WCCP being cleanly shut down.

You cannot shut down an individual WCCP service on a particular port on an WAE; you must shut down WCCP on the WAE. After WCCP is shut down on the WAE, the WAE preserves its WCCP configuration settings.

Configuring Static Bypass Lists for WAEs

Using a static bypass allows traffic flows between a configurable set of clients and servers to bypass handling by the WAE. By configuring static bypass entries on the branch WAE, you can control traffic interception without modifying the router configuration. IP access lists may be configured separately on the router to bypass traffic without first redirecting it to the branch WAE. Typically, the WCCP accept list defines the group of servers that are accelerated (and the servers that are not). Static bypass can be used occasionally when you want to prevent WAAS from accelerating a connection from a specific client to a specific server (or from a specific client to all servers).

Note We recommend that you use ACLs on the WCCP-enabled router where possible, rather than using static bypass lists or interception ACLs on the WAEs, because that is the most efficient method to control traffic interception. If you decide to use static bypass lists or interception ACLs, we recommend using interception ACLs because they are more flexible and give better statistics about passed-through connections. For information about how to configure ACLs on a router, see the "Configuring IP Access Lists on a Router" section on page 4-10. For information about how to configure an interception ACL for a WAE, see the "Configuring an Interception Access Control List" section on page 4-27.

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

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

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

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

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

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

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

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

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

Aggregating Static Bypass Lists

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

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

- Choose **Yes** if you want to configure the device with all static bypass lists that are defined for itself and for device groups to which it belongs.
- Choose **No** if you want to limit the device to just the static bypass lists that are defined for itself.

When you change the setting, you get the following confirmation message: "This option will take effect immediately and will affect the device configuration. Do you wish to continue?" Click **OK** to continue.

Configuring an Interception Access Control List

You can configure an interception ACL to control what incoming traffic across all interfaces is to be intercepted by a WAE device. Packets that are permitted by the ACL are intercepted by the WAE and packets that are denied by the ACL are passed through the WAE without processing.

By configuring an interception ACL on the branch WAE, you can control traffic interception without modifying the router configuration. IP ACLs may be configured separately on the router to bypass traffic without first redirecting it to the branch WAE. Typically, the WCCP accept list defines the group of servers that are accelerated (and the servers that are not). Using an interception ACL allows you to easily bypass uninteresting traffic, for example in a pilot deployment where you do not want to modify the router configuration. Additionally, it allows you to more easily transition from a pilot to a production deployment by allowing and accelerating different kinds of traffic in phases.

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

When used with interface ACLs and WCCP ACLs, the interface ACL is applied first, the WCCP ACL is applied second, and then the interception ACL is applied last. Application policies defined on the WAE are applied after all ACLs have filtered the traffic.

Note The interception ACL feature is mutually exclusive with static bypass lists. You cannot use both types of lists at the same time. We recommend that you use interception ACLs instead of static bypass lists.

To use an interception ACL, first define an ACL (see Chapter 8, "Creating and Managing IP Access Control Lists for WAAS Devices") and then apply it to a device. Interception ACLs are configured for individual devices only and not device groups.

To configure an interception ACL for a WAE, follow these steps:

Step 1 Follow the instructions in Chapter 8, "Creating and Managing IP Access Control Lists for WAAS Devices" to create an ACL that you want to use for interception, but do not apply it to an interface.

Step 2 From the WAAS Central Manager GUI navigation pane, choose **My WAN > Manage Devices**.

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

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

Step 5 Click the arrow control next to the Interception Access List field to display a drop-down list of ACLs you have defined and choose an ACL to apply to interception. Alternatively, you can enter an ACL name directly in the field and create it after you submit this page. If you type in this field, the drop-down list of displayed ACLs is filtered to show only entries beginning with entered text.

If you need to create or edit an ACL, click the **Go to IP ACL** link next to the field to take you to the IP ACL configuration window (this is the Configure > Network > TCP/IP Settings > IP ACL page).

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

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

You can determine if a connection was passed through by an interception ACL by using the **show statistics connection** EXEC command. Flows passed through by an interception ACL are identified with a connection type of "PT Interception ACL."

Additionally, the **show statistics pass-through** command "Interception ACL" counter reports the number of active and completed pass through flows due to an interception ACL.

Configuring Egress Methods for Intercepted Connections

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

- IP forwarding
- WCCP GRE return
- Generic GRE

The default egress method is IP forwarding. If you do not configure the egress method, then the WAE uses IP forwarding. The IP forwarding egress method does not allow you to place WAEs on the same VLAN or subnet as the clients and servers, and it does not ensure that packets are returned to the intercepting router.

The WCCP GRE return and generic GRE egress methods allow you to place WAEs on the same VLAN or subnet as clients and servers. Repeating redirection is prevented by encapsulating the outgoing frames in the GRE frames. Cisco IOS routers handle these GRE frames as bypass frames and do not apply WCCP redirection. With the WCCP GRE return method, WAAS uses the router ID address as the destination for GRE frames; with the generic GRE method, WAAS uses the address of the router configured in the WAE router list.

This technique makes it possible to support redundant routers and router load balancing; WAAS makes a best effort to return frames back to the router from which they arrived, though this is not guaranteed. A notable exception is if flow protection is enabled, the WAE is unable to return flow protected traffic to the originating router because the router information is not available.

Note Designs that require redirected frames to be returned to the originating router are not compatible with the WCCP flow protection feature.

If you want to use this functionality with multiple routers connected to the WAAS network segment, you must ensure connectivity to the router ID address, for example, by configuring static routes. The router ID is the address of the first loopback interface or highest active physical interface. This address can be found in the output of the **show wccp routers** EXEC command.

WAAS applies the following logic in its router selection for WCCP GRE and generic GRE:

- When the WAAS software applies data redundancy elimination (DRE) and compression to a TCP flow, the number of packets that are sent out may be fewer. A single packet that carries optimized data may represent original data that was received in multiple packets redirected from different routers. That optimized data-carrying packet will egress from the WAE to the router that last redirected a packet to the WAE for that flow direction.
- When the WAE receives optimized data, the data may arrive in multiple packets from different routers. The WAAS software expands the optimized data back to the original data, which will be sent out as several packets. Those original data-carrying packets will egress from the WAE to the router that last redirected a packet to the WAE for that flow direction.

The WCCP GRE return and generic GRE egress methods are similar, but the generic GRE egress method is designed specifically to be used in deployments where the router or switch does hardware accelerated processing of GRE packets, such as with the Cisco 7600 series router or the Catalyst 6500 series switch with the Supervisor Engine 32 or 720. Additionally, the generic GRE egress method returns packets to the intercepting router by using a GRE tunnel that you must configure on the router (the WAE end of the tunnel is configured automatically). The generic GRE egress method is supported only when the WCCP GRE interception method is used.

To use the generic GRE egress method, you must create an intercepting router list on the WAE (multicast addresses are not supported) and configure a GRE tunnel interface on each router. For details on configuring GRE tunnel interfaces on the routers, see the "Configuring a GRE Tunnel Interface on a Router" section on page 4-30.

WCCP Version 2 is capable of negotiating the redirect method and the return method for intercepted connections. The WAAS software supports WCCP GRE and WCCP Layer 2 as WCCP negotiated return methods. If WCCP negotiates a WCCP Layer 2 return, the WAE defaults to using IP forwarding as the egress method. The WAE also defaults to IP forwarding if the interception method is set to WCCP Layer 2 and you configure generic GRE as the egress method, which are not compatible. When the WAE defaults to IP forwarding, the WAE logs a minor alarm that is cleared when you correct the configuration so that the interception and egress methods are consistent. The output of the **show egress methods** EXEC command also displays a warning if the interception and egress methods are not consistent.

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

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

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

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

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

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

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

Step 5 Click **Submit**.

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

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

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

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

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

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

The router list must contain the IP address of each intercepting router. Multicast addresses are not supported. Additionally, you must configure a GRE tunnel interface on each router. For details on configuring GRE tunnel interfaces on the routers, see the "Configuring a GRE Tunnel Interface on a Router" section on page 4-30.

To view the egress method that is configured and that is being used on a particular WAE, use the **show egress-methods** EXEC command. To view information about the egress method for each connection segment, use the **show statistics connection egress-methods** EXEC command.

To view the generic GRE tunnel statistics for each intercepting router, use the **show statistics generic-gre** EXEC command. To clear statistics information for the generic GRE egress method, use the **clear statistics generic-gre** EXEC command.

Configuring a GRE Tunnel Interface on a Router

If you plan to use the generic GRE egress method on the WAE, you must configure a GRE tunnel interface on each intercepting router. For ease of configuration, we recommend that you create a single multipoint tunnel on the router, instead of one point-to-point tunnel per WAE in the farm.

If you have only one WAE in the farm, you can use a point-to-point tunnel, however, ensure that the router is configured with no other tunnel that has the same tunnel source as the WAE tunnel.

Note On the Catalyst 6500 series switch with the Supervisor Engine 32 or 720, do not configure more than one GRE tunnel (multipoint or point-to-point) with the same tunnel source interface, otherwise, high switch CPU load can result.

The tunnel interface must have a Layer 3 source interface to which it is attached and this source interface must be the interface whose IP address is configured in the WAE's intercepting router list.

The tunnel interface must be excluded from WCCP interception to avoid routing loops when outbound interception is used. Use the command **ip wccp redirect exclude in**. You can always use this command because it does not cause any impact even when it is not needed, such as for inbound interception.

This section contains the following topics:

Multipoint Tunnel Configuration

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

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

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

The router 1 configuration would look like the following:

```
interface gigabitEthernet 1/1
ip address 192.168.1.1 255.255.255.0
...
interface Tunnel1
ip address 12.12.12.1 255.255.255.0
tunnel source GigabitEthernet1/1
tunnel mode gre multipoint
ip wccp redirect exclude in
end
```

The router 2 configuration would look like the following:

```
interface Vlan815 1/0
ip address 192.168.2.1 255.255.255.0
...
interface Tunnel1
ip address 13.13.13.1 255.255.255.0
tunnel source vlan815
tunnel mode gre multipoint
ip wccp redirect exclude in
end
```

Note Provisioning an IP address for the tunnel interface is what enables it for IP, allowing it to process and forward transit packets. If you do not want to provision an IP address, the tunnel must be IP enabled by making it an IP unnumbered interface. This restricts the tunnel to be a point-to-point tunnel.

Point-To-Point Tunnel Configuration

This section describes how to configure a point-to-point tunnel for a single WAE instead of a multipoint tunnel on the router. A point-to-point tunnel is enabled for IP either by making it unnumbered or by giving it an IP address. The unnumbered method is shown in the following example router configuration:

```
interface gigabitEthernet 1/1
ip address 192.168.1.1 255.255.255.0
...
! Tunnel1 is an unnumbered point-to-point tunnel towards WAE1
interface Tunnel1
ip unnumbered GigabitEthernet1/1
tunnel source GigabitEthernet1/1
! tunnel destination is the IP address of WAE1
tunnel destination 10.10.10.10
```

```
ip wccp redirect exclude in
end
```

Using Policy-Based Routing to Transparently Redirect All TCP Traffic to WAEs

Policy-based routing (PBR), introduced in the Cisco IOS Software Release 11.0, allows you to implement policies that selectively cause packets to take specific paths in the network.

PBR also provides a method to mark packets so that certain kinds of traffic receive differentiated, preferential service when used in combination with queuing techniques enabled through the Cisco IOS software. These queuing techniques provide an extremely powerful, simple, and flexible tool to network managers who implement routing policies in their networks.

PBR enables the router to put packets through a route map before routing them. When configuring PBR, you must create a route map that specifies the match criteria and the resulting action if all of the match clauses are met. You must enable PBR for that route map on a particular interface. All packets arriving on the specified interface matching the match clauses will be subject to PBR.

One interface can have only one route map tag; but you can have several route map entries, each with its own sequence number. Entries are evaluated in order of their sequence numbers until the first match occurs. If no match occurs, packets are routed as usual.

```
Router(config-if)# ip policy route--tag
```

The route map determines which packets are routed next.

You can enable PBR to establish a route that goes through WAAS for some or all packets. WAAS proxy applications receive PBR-redirected traffic in the same manner as WCCP redirected traffic, as follows:

1. In the branch office, define traffic of interest on the branch office router (Edge-Router1) as follows:
   a. Specify which traffic is of interest to the LAN interface (ingress interface) on Edge-Router1.

      Use extended IP access lists to define traffic of interest (traffic from all or filtered local source addresses to any or filtered destination address).

   b. Specify which traffic is of interest to the WAN interface (egress interface) on Edge-Router1.

      Use extended IP access lists to define traffic of interest (traffic from all or filtered local source addresses from any or filtered remote addresses).
2. In the data center, specify which traffic is of interest to the data center router (Core-Router1) as follows:
   a. Specify which traffic is of interest to the LAN interface (ingress interface) on Core-Router1.

      Use extended IP access lists to define traffic of interest (traffic from all or filtered local source addresses to any or filtered destination address).

   b. Specify which traffic is of interest to the WAN interface (egress interface) on Core-Router1.

      Use extended IP access lists to define traffic of interest (traffic from all or filtered local source addresses from any or filtered remote addresses).
3. In the branch office, create route maps on Edge-Router1, as follows:
   a. Create a PBR route map on the LAN interface of Edge-Router1.
   b. Create a PBR route map on the WAN interface of Edge-Router1.

4. In the data center, create route maps on Core-Router1, as follows:
   a. Create a PBR route map on the LAN interface of Core-Router1.
   b. Create a PBR route map on the WAN interface of Core-Router1.
5. In the branch office, apply the PBR route maps to Edge-Router1.
6. In the data center, apply the PBR route maps to Core-Router1.
7. Determine which PBR method to use to verify PBR next-hop availability of a WAE. For more information, see the "Methods of Verifying PBR Next-Hop Availability" section on page 4-38.

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

Figure 4-5 shows that the WAEs (Edge-WAE1 and Core-WAE1) must reside in an out-of-band network that is separate from the traffic's destination and source. For example, Edge-WAE1 is on a subnet separate from the clients (the traffic source), and Core-WAE is on a subnet separate from the file servers and application servers (the traffic destination). Additionally, the WAE may need to be connected to the router that is redirecting traffic to it through a tertiary interface (a separate physical interface) or subinterface to avoid a routing loop. For more information on this topic, see the "Using Tertiary Interfaces or Subinterfaces to Connect WAEs to Routers" section on page 2-23.

Figure 4-5 Example of Using PBR or WCCP Version 2 for Transparent Redirection of All TCP Traffic to WAEs

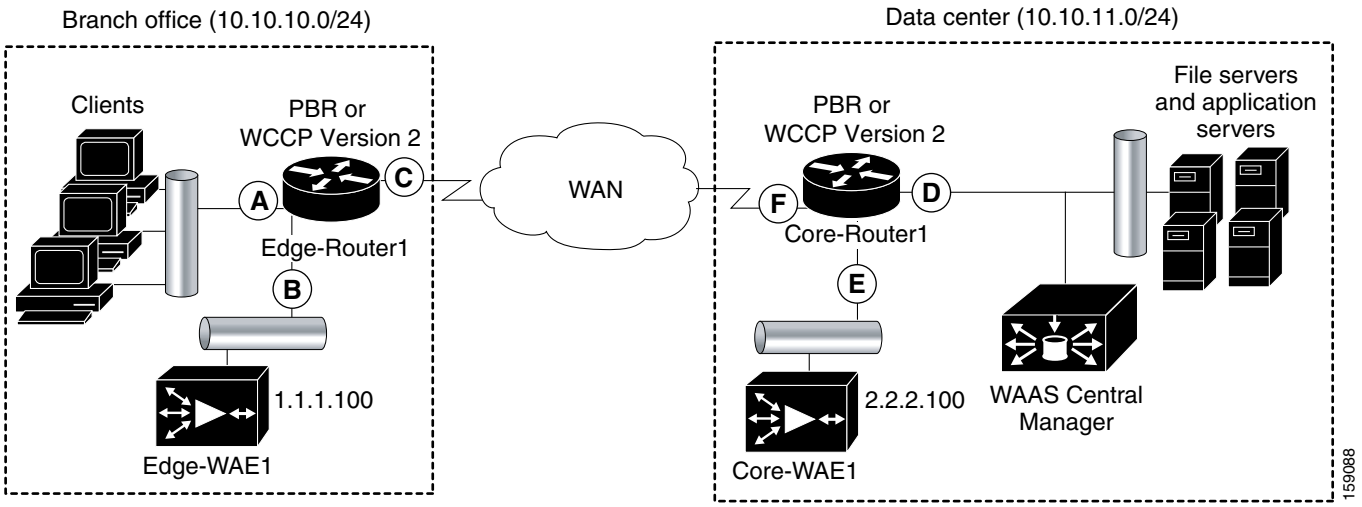

Table 4-3 provides a summary of the router interfaces that you must configure to use PBR or WCCP Version 2 to transparently redirect traffic to a WAE.

Table 4-3 Router Interfaces for WCCP or PBR Traffic Redirection to WAEs

Router interface	Comment
Edge-Router1	
A	Edge LAN interface (ingress interface) that performs redirection on outbound traffic.

Table 4-3 Router Interfaces for WCCP or PBR Traffic Redirection to WAEs (continued)

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

Note In Figure 4-5, redundancy (for example, redundant routers, switches, WAEs, WAAS Central Managers, and routers) is not depicted.

The following example shows how to configure PBR as the traffic redirection method in a WAAS network that has one WAE in a branch office and one WAE in the data center (as shown in Figure 4-5).

Note The commands that are used to configure PBR on a router, can vary based on the Cisco IOS Release installed on the router. For information about the commands that are used to configure PBR for the Cisco IOS Release that you are running on your routers, see the appropriate Cisco IOS configuration guide.

To configure PBR to transparently redirect TCP traffic to WAEs, follow these steps:

Step 1 In the branch office, use extended IP access lists to specify which traffic is of interest to the LAN interface (ingress interface-A) on Edge-Router:

a. On Edge-Router1, define an extended IP access list within the range of 100 to 199. For example, create access list 100 on Edge-Router1:

```
Edge-Router1(config)# ip access-list extended 100
```

b. On Edge-Router1, specify which traffic is of interest to this particular interface:

- For example, mark any IP/TCP traffic from any local source addresses (traffic for any branch office clients) on any TCP port to any destination as interesting:

```
Edge-Router1(config-ext-nac1)# permit tcp 10.10.10.0 0.0.0.255 any
```

- Alternatively, you can selectively mark interesting traffic by defining the source IP subnet, destination IP address, and TCP port numbers. For example, mark IP/TCP traffic from any local source address on TCP ports 135 and 80 to any destination as interesting:

```
Edge-Router1(config-ext-nac1)# permit tcp 10.10.10.0 0.0.0.255 any eq 135
Edge-Router1(config-ext-nac1)# permit tcp 10.10.10.0 0.0.0.255 any eq 80
```

Step 2 In the branch office, use extended IP access lists to specify which traffic is of interest to the WAN interface (egress interface-C) on Edge-Router1:

a. On Edge-Router1, define an extended IP access list within the range of 100 to 199. For example, create access list 101 on Edge-Router1:

```
Edge-Router1(config)# ip access-list extended 101
```

b. On Edge-Router1, specify which traffic is of interest to its WAN interface:

- For example, mark any IP/TCP traffic to a local device as interesting:

```
Edge-Router1(config-ext-nac1)# permit tcp any 10.10.10.0 0.0.0.255
```

- Alternatively, you can selectively mark interesting traffic by defining the source IP subnet, destination IP address, and TCP port numbers. For example, mark IP/TCP traffic to any local source addresses on TCP ports 135 and 80 to any destination as interesting:

```
Edge-Router1(config-ext-nac1)# permit tcp any 10.10.10.0 0.0.0.255 eq 135
Edge-Router1(config-ext-nac1)# permit tcp any 10.10.10.0 0.0.0.255 eq 80
```

Step 3 In the data center, use extended IP access lists to specify which traffic is of interest to the LAN interface (ingress interface-D) on Core-Router1:

a. On Core-Router1, define an extended IP access list within the range of 100 to 199. For example, create access list 102 on Core-Router1:

```
Core-Router1(config)# ip access-list extended 102
```

b. On Core-Router1, specify which traffic is of interest to its LAN interface:

- For example, mark any IP/TCP traffic sourced from any local device (for example, traffic sourced from any file server or application server in the data center) on any TCP port to any destination as interesting:

```
Core-Router1(config-ext-nac1)# permit tcp 10.10.11.0 0.0.0.255 any
```

- Alternatively, you can selectively mark traffic as interesting by defining the source IP subnet, destination IP address, and TCP port numbers. For example, selectively mark IP/TCP traffic sourced from any local device on TCP ports 135 and 80 to any destination as interesting:

```
Core-Router1(config-ext-nac1)# permit tcp 10.10.11.0 0.0.0.255 any eq 135
Core-Router1(config-ext-nac1)# permit tcp 10.10.11.0 0.0.0.255 any eq 80
```

Step 4 In the data center, use extended IP access lists to mark traffic of interest for the WAN interface (egress interface-F) on Core-Router1:

a. On Core-Router1, define an extended access list within the range of 100 to 199. For example, create access list 103 on Core-Router1:

```
Core-Router1(config)# ip access-list extended 103
```

b. On Core-Router1, mark interesting traffic for the WAN interface:

- For example, mark any IP/TCP traffic destined to any local device (for example, traffic destined to any file server or application server in the data center) as interesting:

```
Core-Router1(config-ext-nac1)# permit tcp any 10.10.11.0 0.0.0.255
```

- Alternatively, you can selectively mark traffic as interesting by defining the source IP subnet, destination IP address, and TCP port numbers. For example, mark IP/TCP traffic on ports 135 and 80 to any local source addresses as interesting:

```
Core-Router1(config-ext-nac1)# permit tcp any 10.10.11.0 0.0.0.255 eq 135
Core-Router1(config-ext-nac1)# permit tcp any 10.10.11.0 0.0.0.255 eq 80
```

Step 5 In the branch office, define PBR route maps on Edge-Router1:

a. Define a route map for the LAN interface (ingress interface). In the following example, the WAAS-EDGE-LAN route map is created:

```
Edge-Router1(config)# route-map WAAS-EDGE-LAN permit
```

b. Define a route map for the WAN interface (egress interface).

In the following example, the WAAS-EDGE-WAN route map is created:

```
Edge-Router1(config)# route-map WAAS-EDGE-WAN permit
```

c. Specify the match criteria.

Use the **match** command to specify the extended IP access list that Edge-Router1 should use to determine which traffic is of interest to its WAN interface. If you do not specify a **match** command, the route map applies to all packets.

In the following example, Edge-Router1 is configured to use the access list 101 as the criteria for determining which traffic is of interest to its WAN interface:

```
Edge-Router1(config-route-map)# match ip address 101
```

Note The **ip address** command option matches the source or destination IP address that is permitted by one or more standard or extended access lists.

d. Specify how the matched traffic should be handled.

In the following example, Edge-Router1 is configured to send the packets that match the specified criteria to the next hop, which is Edge-WAE1 that has an IP address of 1.1.1.100:

```
Edge-Router1(config-route-map)# set ip next-hop 1.1.1.100
```

Note If you have more than one branch WAE, you can specify the IP address of a second branch WAE for failover purposes (for example, enter the **set ip next-hop 1.1.1.101** command on Edge-Router1) to specify a next-hop address of 1.1.1.101 (the IP address of Edge-WAE2) for failover purposes. The **next-hop** command is used for failover purposes and not for load-balancing purposes.

Step 6 In the data center, create route maps on Core-Router1:

a. Define a route map on the LAN interface (ingress interface).

In the following example, the WAAS-CORE-LAN route map is created:

```
Core-Router1(config)# route-map WAAS-CORE-LAN permit
```

b. Define a route map on the WAN interface (egress interface).

In the following example, the WAAS-CORE-WAN route map is created:

```
Core-Router1(config)# route-map WAAS-CORE-WAN permit
```

c. Specify the match criteria.

Use the **match** command to specify the extended IP access list that Core-Router 1 should use to determine which traffic is of interest to its WAN interface. If you do not enter a **match** command, the route map applies to all packets. In the following example, Core-Router1 is configured to use the access list 103 as the criteria for determining which traffic is of interest to its WAN interface:

```
Core-Router1(config-route-map)# match ip address 103
```

d. Specify how the matched traffic is to be handled.

In the following example, Core-Router1 is configured to send packets that match the specified criteria to the next hop, which is Core-WAE1 that has an IP address of 2.2.2.100:

```
Core-Router1(config-route-map)# set ip next-hop 2.2.2.100
```


Note If you have more than one data center WAE, you can specify the IP address of a second data center WAE for failover purposes (for example, enter the **set ip next-hop 2.2.2.101** command on Core-Router1) to specify a next-hop address of 2.2.2.101 (the IP address of Core-WAE2) for failover purposes. The **next-hop** command is used for failover purposes and not for load-balancing purposes.

Step 7 In the branch office, apply the route maps to the LAN interface (ingress interface) and the WAN interface (egress interface) on Edge-Router1:

a. On Edge-Router1, enter interface configuration mode:

```
Edge-Router1(config)# interface FastEthernet0/0.10
```

b. Specify that the LAN router interface should use the WAAS-EDGE-LAN route map for PBR:

```
Edge-Router1(config-if)# ip policy route-map WAAS-EDGE-LAN
```

c. Enter interface configuration mode:

```
Edge-Router1(config-if)# interface Serial0
```

d. Specify that the WAN router interface should use the WAAS-EDGE-WAN route map for PBR:

```
Edge-Router1(config-if)# ip policy route-map WAAS-EDGE-WAN
```

Step 8 In the data center, apply the route maps to the LAN interface (ingress interface) and the WAN interface (egress interface) on Core-Router1:

a. On Core-Router1, enter interface configuration mode:

```
Core-Router1(config)# interface FastEthernet0/0.10
```

b. Specify that for PBR, the LAN router interface should use the WAAS-CORE-LAN route map:

```
Core-Router1(config-if)# ip policy route-map WAAS-CORE-LAN
```

c. Enter interface configuration mode:

```
Core-Router1(config-if)# interface Serial0
```

d. Specify that for PBR, the WAN router interface should use the WAAS-CORE-WAN route map:

```
Core-Router1(config-if)# ip policy route-map WAAS-CORE-WAN
```

Methods of Verifying PBR Next-Hop Availability

When using PBR to transparently redirect traffic to WAEs, we recommend that you use one of the following methods to verify the PBR next-hop availability of a WAE. The method that you choose is based on the version of the Cisco IOS software that is running on the routers and the placement of your WAEs. However, method 2 is the preferred method whenever possible:

- Method 1—If the device sees the WAEs as a CDP neighbor (directly connected), it can use CDP and ICMP to verify that the WAE is operational. For more information, see the "Method 1: Using CDP to Verify Operability of WAEs" section on page 4-38.
- Method 2 (Recommended method)—If the device is running the Cisco IOS software Release 12.4 or later and the device does not see the WAE as a CDP neighbor, IP service level agreements (SLAs) can be used to verify that the WAE is operational using ICMP echoes. For more information, see the "Method 2: Using IP SLAs to Verify WAE Operability Using ICMP Echo Verification (Recommended Method)" section on page 4-39.
- Method 3—If the device is running the Cisco IOS software Release 12.4 or later and does not see the WAE as a CDP neighbor, IP SLAs can be used to verify that the WAE is operational using TCP connection attempts. For more information, see the "Method 3: Using IP SLAs to Verify WAE Operability Using TCP Connection Attempts" section on page 4-40.

Note In this section, the term device is used to refer to the router or switch that has been configured to use PBR to transparently redirect traffic to a WAE.

To verify whether the WAE is CDP visible to a device that has been configured to use PBR, enter the **show cdp neighbors** command on the device. If the WAE is CDP visible to the device, the WAE will be listed in the output of the **show cdp neighbors** command.

Method 1: Using CDP to Verify Operability of WAEs

If the device that is configured to use PBR views the WAEs as a CDP neighbor (the WAE is directly connected to the device), you can configure CDP and ICMP to verify the availability of a WAE as a PBR next hop.

The following example shows how to use this method to verify PBR next-hop availability of a WAE. You must complete the following configuration process for each of the LAN and WAN route maps that are configured when CDP should be used.

To use CDP to verify operability of WAEs, follow these steps:

Step 1 On the router where PBR is configured (for example, on the branch office router named Edge-Router1), enter configuration mode and enable CDP on the router:

```
Edge-Router1(config)# cdp run
```

Step 2 Enable route-map configuration mode for the route map, WAAS-EGDE-LAN, which has already been created on the router:

```
Edge-Router1(config)# route-map WAAS-EDGE-LAN permit
```

Step 3 Configure the router to use CDP to verify the availability of the configured next-hop addresses:

```
Edge-Router1(config-route-map)# set ip next-hop verify-availability
```

Step 4 Enable CDP on the WAE (for example, on the branch office WAE named Edge-WAE1) that you want the router to redirect traffic to using PBR:

```
Edge-WAE1(config)# cdp enable
```

If you are configuring PBR and have multiple WAEs and are using Method 1 to verify the PBR next-hop availability of a WAE, no additional configuration is necessary after you have completed the preceding process.

Method 2: Using IP SLAs to Verify WAE Operability Using ICMP Echo Verification (Recommended Method)

To use IP SLAs and ICMP (the recommended method) to verify PBR next-hop availability of a WAE, follow these steps:

Step 1 On the branch office router named Edge-Router1, enter the route-map configuration mode for the route map named WAAS-EDGE-LAN, which has been previously configured on this router:

```
Edge-Router1(config)# route-map WAAS-EDGE-LAN permit
```

Step 2 Specify a match condition for the traffic. In the following example, the match condition specifies access list number 105:

```
Edge-Router1(config)# match ip address 105
```

Step 3 Configure the route map to use IP SLA tracking instance number 1 to verify the availability of the next-hop WAE (for example, the branch WAE named Edge-WAE1 that has an IP address of 1.1.1.100):

```
Edge-Router1(config-route-map)# set ip next-hop verify-availability 1.1.1.100 track 1
```

Note Enter the **set ip next-hop verify-availability** command for each route-map that has been configured on this branch office edge router and on the data center's core router that has also been configured to use PBR to redirect traffic to WAEs.

Step 4 Configure the IP SLA tracking instance 1:

```
Edge-Router1(config-route-map)# exit
Edge-Router1(config)# ip sla 1
Edge-Router1(config-ip-sla)#
```

Step 5 Configure the router to echo Edge-WAE1 using the specified source interface:

```
Edge-Router1(config-ip-sla)# icmp-echo 1.1.1.100 source-interface FastEthernet 0/0.20
```

Step 6 Configure the router to perform the echo every 20 seconds:

```
Edge-Router1(config-ip-sla)# frequency 20
Edge-Router1(config-ip-sla)# exit
```

Step 7 Schedule the IP SLA tracking instance 1 to start immediately and to run continuously:

```
Edge-Router1(config)# ip sla schedule 1 life forever start-time now
```

Step 8 Configure the IP SLA tracking instance 1 to track the device, which is defined in the IP SLA tracking instance 1:

```
Edge-Router1(config)# track 1 rtr 1
```

If you are configuring PBR and have multiple WAEs, and you are using Method 2 to verify PBR next-hop availability of a WAE, you must configure a separate IP SLA per WAE and then run the **track** command per IP SLA.

Method 3: Using IP SLAs to Verify WAE Operability Using TCP Connection Attempts

If the device that is configured for PBR is running the Cisco IOS software Release 12.4 or later and does not see the WAE as a CDP neighbor, IP SLAs can be used to verify that the WAE is alive using TCP connection attempts. IP SLAs can be used to monitor a WAE's availability as the PBR next hop using TCP connection attempts at a fixed interval of 60 seconds.

To verify PBR next-hop availability of a WAE, follow these steps:

Step 1 On the branch office router named Edge-Router1, enter route-map configuration mode for the route map named WAAS-EDGE-LAN, which has been previously configured on this router:

```
Edge-Router1(config)# route-map WAAS-EDGE-LAN permit
```

Step 2 Configure the route map to use IP SLA tracking instance number 1 to verify the availability of the next-hop WAE (the Edge WAE that has an IP address of 1.1.1.100):

```
Edge-Router1(config-route-map)# set ip next-hop verify-availability 1.1.1.100 track 1
```

Note Enter the **set ip next-hop verify-availability** command for each route map that is configured on this branch office edge router and on the data center's core router that has also been configured to use PBR to transparently redirect traffic to WAEs.

Step 3 Configure the IP SLA tracking instance 1:

```
Edge-Router1(config-route-map)# exit
Edge-Router1(config)# ip sla 1
```

Step 4 Configure the router to use the specified source and destination ports to use TCP connection attempts at a fixed interval of 60 seconds to monitor the WAE availability:

```
Edge-Router1(config-ip-sla)# tcp-connect 1.1.1.100 80 source-port 51883 control disable
Edge-Router1(config-ip-sla)# exit
```

Step 5 Schedule the IP SLA tracking instance 1 to start immediately and to run forever:

```
Edge-Router1(config)# ip sla schedule 1 life forever start-time now
```

Step 6 Configure the IP SLA tracking instance 1 to track the device, which is defined in the IP SLA tracking instance 1:

```
Edge-Router1(config)# track 1 rtr 1
```

If you are configuring PBR and have multiple WAEs, and you are using Method 3 to verify PBR next-hop availability of a WAE, you must configure a separate IP SLA per WAE and then run the **track** command per IP SLA.

Using Inline Mode to Transparently Intercept TCP Traffic

The WAE can physically and transparently intercept traffic between the clients and the router by using inline mode. To use inline mode, you must use a WAE with the Cisco WAE Inline Network Adapter or Interface Module installed. In this mode, you physically position the WAE device in the path of the traffic that you want to optimize, typically between a switch and a router, as shown in Figure 4-6. Redirection of traffic is not necessary.

Note When you install any inline WAE device, you must follow the cabling requirements described in the "Cabling" section of *Installing the Cisco WAE Inline Network Adapter* located on Cisco.com.

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

Figure 4-6 Inline Interception

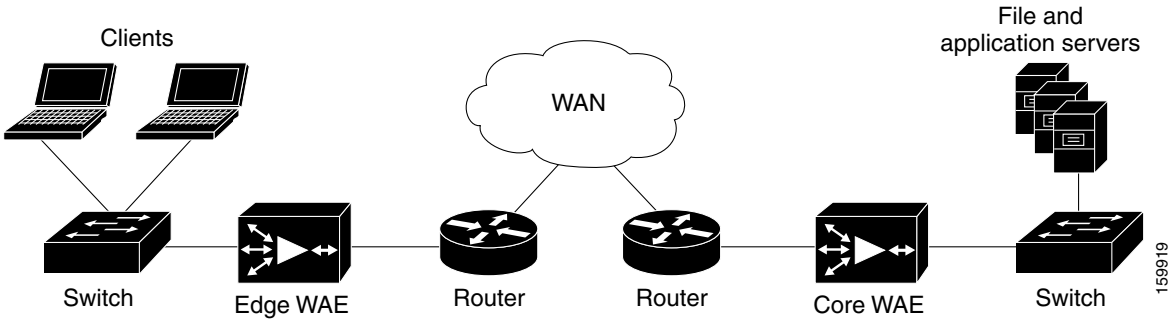

Note Inline mode and WCCP redirection are exclusive. You cannot configure inline mode if the WAE is configured for WCCP operation. Inline mode is the default mode when a Cisco WAE Inline Network Adapter is installed in a WAE device, but you must configure inline mode explicitly on a device with a Cisco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

Note An inline WAE can be configured as a Central Manager, but the inline interception functionality is not be available.

The Cisco WAE Inline Network Adapter contains two or four Ethernet ports and the Cisco Interface Module contains four to eight Ethernet ports. Ports on the Cisco WAE Inline Network Adapter are always configured as inline ports, while ports on the Cisco Interface Module are configured as normal standalone ports by default and you must explicitly configure these ports as inline ports. Each pair of inline ports is grouped into a logical inline group.

Each inline group has one LAN-facing port and one WAN-facing port. Typically, you use just one inline group, and connect the LAN-facing port to a switch and the WAN-facing port to a router. On adapters or interface modules with additional ports, the additional groups of interfaces are provided if you are using a network topology where you need to connect the WAE to multiple routers. Traffic that enters on one interface in a group exits the device on another interface in the same group.

Hardware platform support for inline ports is as follows:

- WAE-512/612—Support one installed four-port Cisco WAE Inline Network Adapter.
- WAVE-274/474—Support one installed two-port Cisco WAE Inline Network Adapter.
- WAVE-574—Supports one installed two-port or four-port Cisco WAE Inline Network Adapter.
- WAE-674/7341/7371—Support up to two installed four-port Cisco WAE Inline Network Adapters, providing a total of eight inline ports.
- WAVE-294/594/694/7541/7571/8541—Support one installed Cisco Interface Module, which can be configured with four or eight ports.

Note The 10-Gigabit Cisco Interface Module cannot be used in inline mode.

You have the option of assigning an IP address to an inline interface, but it is not required. For more information, see the "Configuring an IP Address on an Inline Interface" section on page 4-47.

Traffic that flows through an inline group is transparently intercepted for optimization. Traffic that does not need to be optimized is bridged across the LAN/WAN interfaces. If a power, hardware, or unrecoverable software failure occurs, the network adapter automatically begins operating in bypass mode (fail-close), where all traffic is mechanically bridged between the LAN and WAN interfaces in each group. The Cisco WAE Inline Network Adapter and Cisco Interface Module also operate in bypass mode when the WAE is powered off or starting up. Additionally, you can manually put an inline group into bypass mode.

Inline mode is configured by default to accept all TCP traffic. If the network segment in which the WAE is inserted is carrying 802.1Q tagged (VLAN) traffic, initially traffic on all VLANs is accepted. Inline interception can be enabled or disabled for each VLAN. However, optimization policies cannot be customized based on the VLAN.

You can serially cluster WAE devices operating in inline mode to provide higher availability if a device fails. For details, see the "Clustering Inline WAEs" section on page 4-49.

Note When a WAE inline group enters bypass mode, the switch and router ports to which it is connected may have to reinitialize, which may cause an interruption of several seconds in the traffic flow through the WAE.

If the WAE is deployed in a configuration where the creation of a loop is not possible (that is, if it is deployed in a standard fashion between a switch and a router), configure PortFast on the switch port to which the WAE is connected. PortFast allows the port to skip the first few stages of the Spanning Tree Algorithm (STA) and move more quickly into a packet forwarding mode.

This section contains the following topics:

Enabling Inline Operation

On WAVE-294/594/694/7541/7571/8541 devices that use Cisco Interface Modules, the Interface Module ports are configured by default for normal standalone operation. If you want to use the device in inline mode, you must configure the ports for inline operation. Enabling inline mode configures all ports for inline operation.

On other WAE devices, which use the Cisco WAE Inline Network Adapter, this operation is unnecessary because the ports on the adapter always operate in inline mode. However, you can use this configuration window to enable or disable VLAN ID connection checking, which is the only setting that appears for such WAE devices.

To enable inline operation and configure general settings, follow these steps:

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

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

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

The Device Dashboard window appears.

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

Step 4 The General Settings window appears. (See Figure 4-7.)

Figure 4-7 Inline General Settings Window

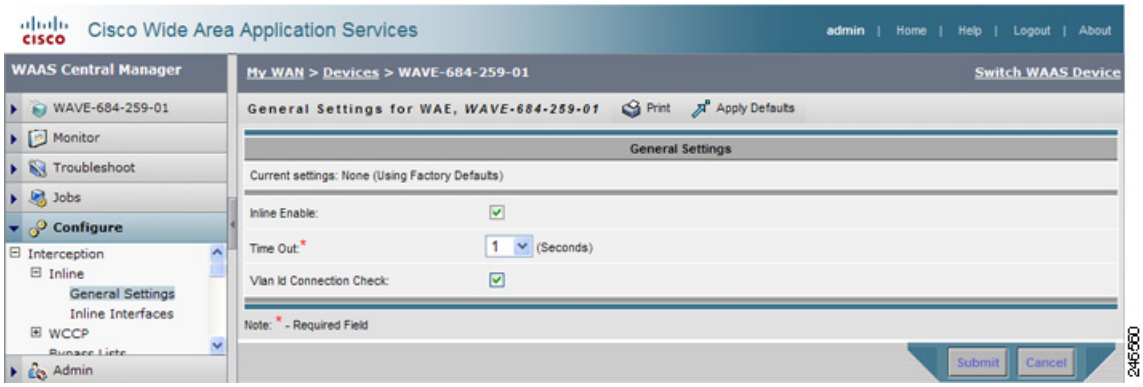

Step 5 Check the **Inline Enable** check box to enable inline operation.

This check box appears only for WAVE devices that use Cisco Interface Modules. For devices that use Cisco WAE Inline Network Adapters, the failover timeout is configured in the Inline Interface Settings window (see Figure 4-8 on page 4-45).

Step 6 From the Time Out drop-down list, choose the failover timeout (1, 5 or 25 seconds), which is the number of seconds that the interface should wait before going into bypass mode, after a device or power failure. The default is 1 second. (This check box appears only for WAVE devices that use Cisco Interface Modules.)

Step 7 Check the **VLAN ID Connection Check** check box to enable VLAN ID connection checking. Uncheck the check box to disable it. The default setting is enabled.

WAAS uses the VLAN ID to intercept or bridge VLAN traffic on the inline interface for a TCP flow. The VLAN ID of all packets sent in a particular TCP connection must match; any packets with a different VLAN ID will be bridged and not optimized. If your system has an asymmetric routing topology, in which the traffic flow in one direction uses a different VLAN ID than the traffic flow from the other direction, you may need to disable VLAN ID checking to ensure that the traffic is optimized.

Step 8 Click **Submit**.

After enabling inline mode, it takes about two datafeed poll cycles (about 10 minutes by default) for the inline groups to appear in the Inline Interfaces list.

Note If you configure any of the interfaces on a Cisco Interface Module with nondefault settings (standby group, port channel, BVI, speed, duplex, IP address, ACLs, and so on), inline mode cannot be enabled and a warning message appears that tells you to check all interfaces for any configuration settings. You must remove all configuration settings from all interface module interfaces (slot 1) and then return to this configuration window to enable inline mode.

To enable inline operation and set the failover timeout from the CLI, use the **inline enable** global configuration command.

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

Configuring Inline Interface Settings

To configure inline interface settings, follow these steps:

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

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

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

The Device Dashboard window appears.

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

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

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

Figure 4-8 Modifying Inline Interface Window

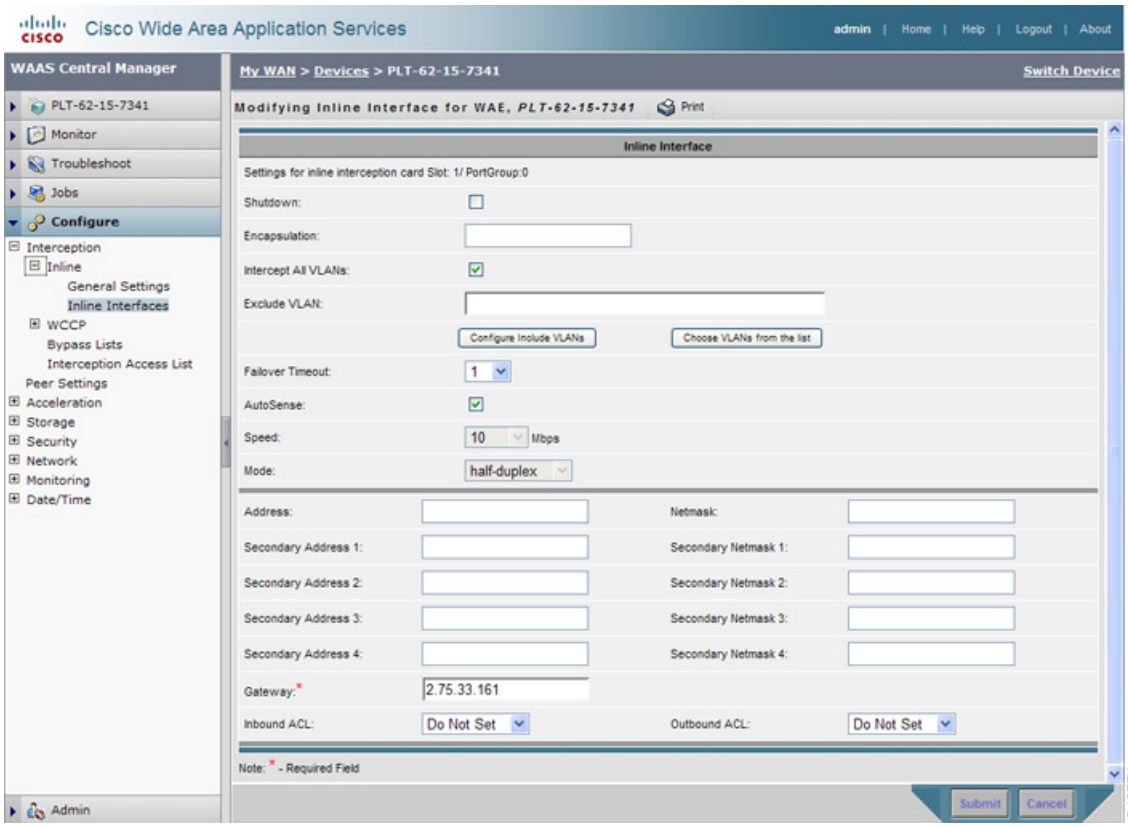

Step 4 Check the **Shutdown** check box to shut down the inline group. This setting bridges traffic across the LAN/WAN interfaces without any processing.

Step 5 In the Encapsulation field, enter the VLAN ID that is to be assigned to traffic that leaves the WAE. The VLAN ID should be set to match the VLAN ID expected by the router.

For more information about the VLAN ID, see the "Configuring an IP Address on an Inline Interface" section on page 4-47.

Step 6 Check the **Intercept all VLANs** check box to enable inline interception on the interface group. Inline interception is enabled by default when the WAE contains a Cisco WAE Inline Network Adapter but must be explicitly enabled on devices with a Cisco Interface Module (see the "Enabling Inline Operation" section on page 4-43).

Note Inline mode and WCCP redirection are exclusive. You cannot configure inline mode if the WAE is configured for WCCP operation. Inline mode is the default mode when the Cisco WAE Inline Network Adapter is installed in a WAE device.

Step 7 In the Exclude VLAN field, enter a list of one or more VLAN ranges to exclude from optimization. You can enter the word "native" to exclude the native VLAN. Separate each VLAN range from the next with a comma. Alternatively, you can select VLAN ranges from a list by following these steps:

a. Click the **Configure Include VLANs** button when you know the list of VLANs that you want to include in inline interception. This button runs a script that prompts you for a comma-separated list of VLANs that you want to include. The script generates an inverse list of all VLANs that should be excluded and then updates the window and puts the list into the Exclude VLAN field.

b. Click the **Choose VLANs from the list** button to pick VLAN ranges. The VLAN Range Assignments window appears, displaying the VLAN ranges that are defined. Defining VLAN ranges is described in the "Configuring VLANs for Inline Support" section on page 4-48.

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

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

d. Click **Submit**.

Step 8 From the Failover Timeout drop-down list, choose **1**, **3**, **5**, or **10** seconds. The default is 1 second. This value sets the number of seconds after a failure event that the WAE waits before beginning to operate in bypass mode. In bypass mode, all traffic received on either port of the interface group is forwarded out the other port in the group.

This check box applies only to devices that use Cisco WAE Inline Network Adapters. For devices that use Cisco Interface Modules, the failover timeout is configured in the Inline General Settings window (see Figure 4-7 on page 4-43) and does not appear in this window.

Step 9 Configure the Speed and Mode port settings as follows (these settings are not used on 10-Gigabit Ethernet interfaces):

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

b. From the Speed drop-down list, choose a transmission speed (**10**, **100**, or **1000** Mbps). You must choose 1000 Mbps for fiber Gigabit Ethernet interfaces on a Cisco Interface Module.

c. From the Mode drop-down list, choose a transmission mode (**full-duplex** or **half-duplex**). You must choose full-duplex for fiber Gigabit Ethernet interfaces on a Cisco Interface Module.

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

Step 10 In the Address field, enter an IP address for the inline interface, if you want to assign an IP address.

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

Step 12 Enter up to four secondary IP addresses and corresponding subnet masks in the Secondary Address and Secondary Netmask fields.

Configuring multiple IP addresses allows the device to be present in more than one subnet and can be used to optimize response time because it allows the data to go directly from the WAAS device to the client that is requesting the information without being redirected through a router. The WAAS device becomes visible to the client because both are configured on the same subnet.

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

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

The drop-down list contains all the IP ACLs that you configured in the system.

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

Step 16 Click **Submit**.

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

Configuring an IP Address on an Inline Interface

You can assign IP addresses to the inline group interfaces but it is not required. You can assign a primary IP address and up to four secondary IP addresses, using the procedure discussed in the "Configuring Inline Interface Settings" section on page 4-44.

You can set an inline group interface as the primary interface on the WAE by using the Configure > Network > Network Interfaces window, in the Primary Interface drop-down list.

In scenarios where the primary interface for a WAE is set to an inline group interface and management traffic is configured on a separate IP address (either on a secondary IP address on the same inline group interface or on a built-in interface), you must configure the WAAS Central Manager to communicate with the WAE on the IP address designated for management traffic. Configure the WAE management traffic IP address in the *Device Name* > Activation window, in the Management IP field.

If a WAE operating in inline mode is present in an 802.1Q VLAN trunk line between a switch and a router, and you are configuring the inline interface with an IP address, you must set the VLAN ID that is to be assigned to traffic that leaves the WAE. The VLAN ID should be set to match the VLAN ID expected by the router.

Use the **encapsulation dot1Q** interface command to assign a VLAN ID, as follows:

```
(config)# interface inlineGroup 1/0
(config-if)# encapsulation dot1Q 100
```

This example shows how to assign VLAN ID 100 to the traffic leaving the WAE. The VLAN ID can range from 1 through 4094.

Note You can set the VLAN ID of the inline traffic by using the **encapsulation dot1Q** interface command or by using the Central Manager page at **Configure > Interception > Inline > Inline Interfaces** (see the "Configuring Inline Interface Settings" section on page 4-44).

If the VLAN ID that you set does not match the VLAN ID expected by the router subinterface, you may not be able to connect to the inline interface IP address.

The inline adapter supports only a single VLAN ID for each inline group interface. If you have configured a secondary address from a different subnet on an inline interface, you must have the same secondary address assigned on the router subinterface for the VLAN.

Using IEEE 802.1Q tunneling increases the frame size by 4 bytes when the tag is added. Therefore, you must configure all switches through which the tunneled packet traverses to be able to process larger frames by increasing the device MTU to at least 1504 bytes.

The following operating considerations apply to configuring IP addresses on the inline interfaces:

- This feature provides basic routable interface support and does not support the following additional features associated with the built-in interfaces: standby and port channel.

- If you have configured a WAE to use the inline interfaces for all traffic, inline interception must be enabled or the WAE will not receive any traffic.
- If you have configured a WAE to use the inline interfaces for all traffic and it goes into mechanical bypass mode, the WAE become inaccessible through the inline interface IP address. Console access is required for device management when an inline interface is in bypass mode.
- If you have configured a WAE with an IP address on an inline interface, the interface can accept only traffic addressed to it and ARP broadcasts, and the interface cannot accept multicast traffic.
- In a deployment using the Hot Standby Router Protocol (HSRP) where two routers that participate in an HSRP group are directly connected through two inline groups, HSRP works for all clients if the active router fails. However, this redundancy does not apply to the IP address of the WAE itself for management traffic, if management traffic is also configured to use the inline interface. If the active router fails, you will not be able to connect to the WAE inline IP address because the inline interface is physically connected to the failed router interface. You will be able to connect to the WAE through the second inline group interface that is connected to the standby router. If redundancy is needed for the IP address of the WAE itself for management traffic, we recommend that you use the IP addresses of the built-in interfaces rather than the inline interfaces.

Configuring VLANs for Inline Support

Initially, the WAE accepts traffic from all VLANs. You can configure the WAE to include or exclude traffic from certain VLANs; for excluded VLANs, traffic is bridged across the LAN/WAN interfaces in a group and is not processed.

To configure a VLAN for inline support, follow these steps:

Step 1 From the WAAS Central Manager GUI navigation pane, choose **Configure > Platform > Vlans**.

The Vlans window appears, listing the VLANs that are defined. You can click the **Edit Vlan** icon next to an existing VLAN that you want to modify.

Step 2 In the taskbar, click the **Create New Vlan** icon. The Creating VLAN window appears. (See Figure 4-9.)

Figure 4-9 Creating VLAN Window

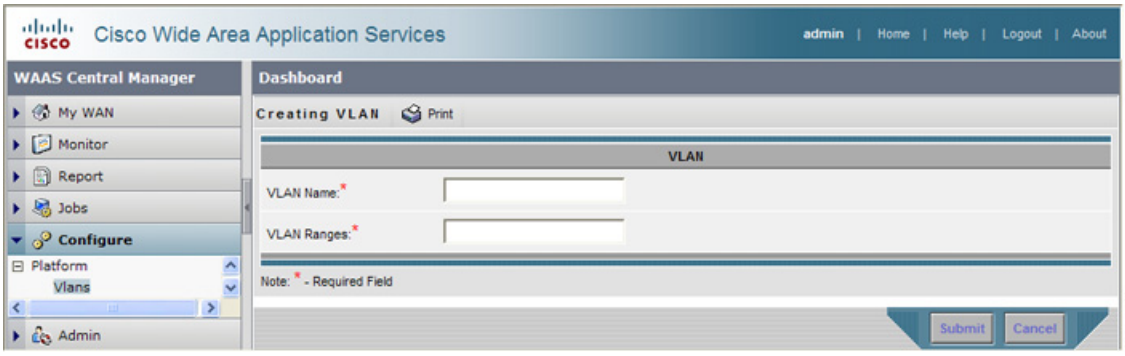

Step 3 In the VLAN Name field, enter a name for the VLAN list.

Step 4 In the VLAN Ranges field, enter a list of one or more VLAN ranges. Separate each VLAN range from the next with a comma (but no space). This list of VLAN ranges can be included or excluded from optimization when you configure the inline interface group, as described in the "Configuring Inline Interface Settings" section on page 4-44. You cannot specify the term "native" in this field.

Step 5 Click **Submit**.

This facility for creating VLAN lists is provided so that you can configure VLAN lists globally. You do not need to use this facility to configure VLANs for an inline interface. You can configure VLANs directly in the inline interface settings window, as described in the "Configuring Inline Interface Settings" section on page 4-44.

Clustering Inline WAEs

You can serially cluster two WAE devices that are operating in inline mode to provide higher availability in the data center if a device fails. If the current optimizing device fails, the inline group shuts down, or the device becomes the overloaded, the second WAE device in the cluster provides the optimization services. Deploying WAE devices in a serial inline cluster for scaling or load balancing is not supported.

Note Overload failover occurs on TFO overload, not overload of individual application accelerators, and it is intended for temporary overload protection. We do not recommend that you continually run a WAE in an overloaded state, frequently triggering overload failover.

A serial cluster consists of two WAE devices connected together sequentially in the traffic path. The WAN port of one device is connected to the LAN port of the next device, as shown in Figure 4-10.

Figure 4-10 Inline Cluster

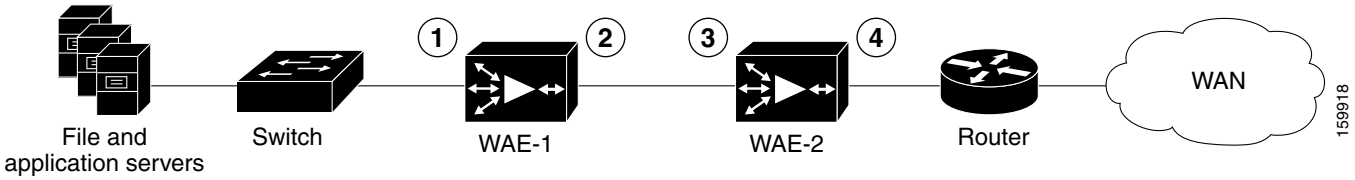

1	Inline LAN port on WAE-1	**3**	Inline LAN port on WAE-2
2	Inline WAN port on WAE-1	**4**	Inline WAN port on WAE-2

In a serial cluster, all traffic between the switch and router passes through all inline WAEs. In Figure 4-10, TCP connections are optimized by WAE-1. If WAE-1 fails, it bypasses the traffic and connections are then optimized by WAE-2.

The policy configuration of serially clustered WAEs should be the same. Additionally, we recommend that you use the same device for both WAEs in the cluster.

When serially clustering inline WAEs, on each WAE you must configure the address of the other WAE in the cluster as a non-optimizing peer. This disables optimization between the two peer WAEs in the serial cluster, since you want optimization only between the WAE peers on each side of the WAN link.

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

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

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

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

The Device Dashboard window appears.

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

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

Figure 4-11 Peer Settings Window

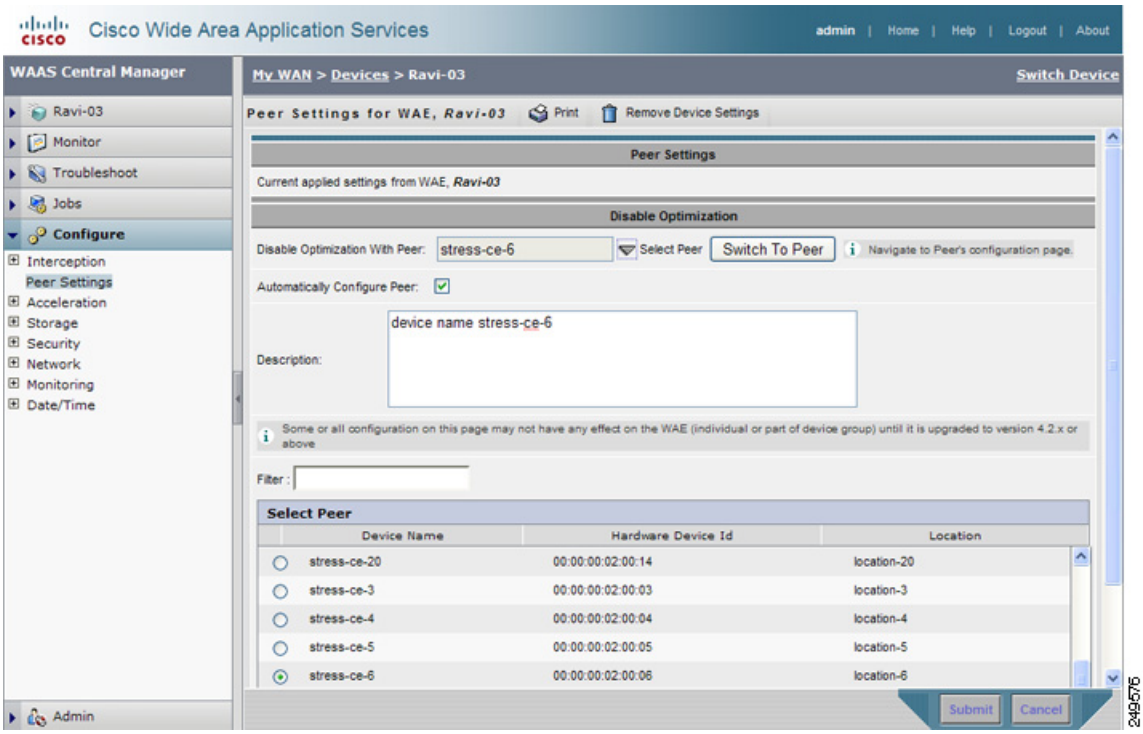

Step 4 Click the Select Peer triangle control to display in the lower part of the window other WAEs that are registered with this Central Manager (see the Select Peer area).

Step 5 In the Select Peer area, click the radio button next to the serial peer of the current device. The peer device name appears in the Disable Optimization With Peer field.

If you need to filter the device list, enter a string in the Filter field. As you enter characters, the device list is dynamically filtered to include only devices that have the filter string in their name or hardware device ID.

Step 6 Check the Automatically Configure Peer check box to allow the Central Manager to configure the other peer with a similar setting to disable optimization with the current device.

If you do not check this box, you must manually configure the other peer to disable optimization with the current device. After you submit your changes, you can click the **Switch to Peer** button to go to this same configuration page for the peer device.

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

Step 8 Click **Submit**.

To disable serial peer optimization from the CLI, use the **no peer device-id** global configuration command. To reenable serial peer optimization, use the **peer device-id** global configuration command.

To view the status of all serial cluster pairs registered with the Central Manager, from the WAAS Central Manager GUI navigation pane, choose **Configure > Peer Settings**. The Peer Settings status window appears, as shown in Figure 4-12.

Figure 4-12 Peer Settings For All Devices Window

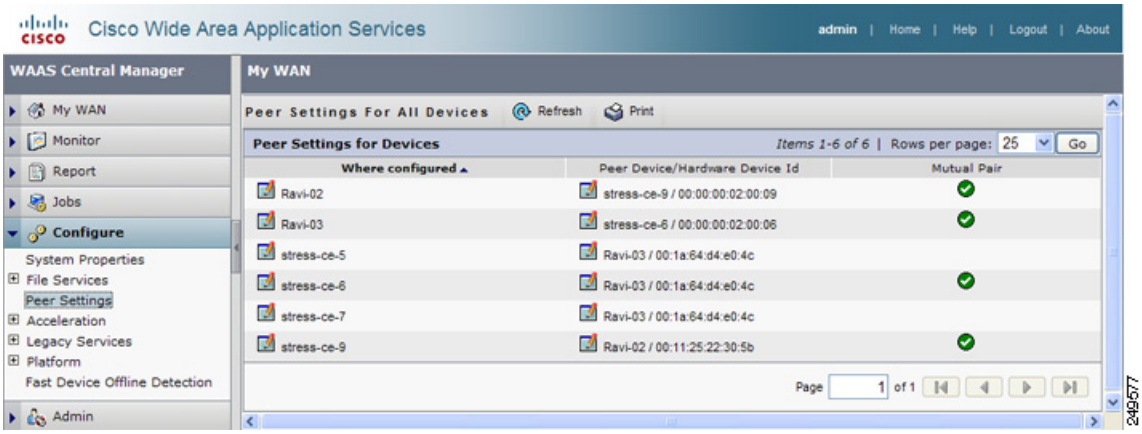

The window lists each WAE for which you have configured peer optimization settings. Verify that there are two entries for each serial cluster pair, both with a check mark in the Mutual Pair column. There should be an entry for each WAE in the pair (for example, the first and last entries in the figure).

If you see an entry without a check mark in the Mutual Pair column (like the third one in the figure), it indicates a WAE on which a serial peer is configured, but the peer is not similarly configured with the first device as its serial peer.

Configuring VPATH Interception on a vWAAS Device

VPATH intercepts traffic from the VM server, redirects it to a vWAAS device for WAN optimization, and then returns the response back to the Virtual Ethernet Module (VEM). The vWAAS egress traffic received by the VEM is forwarded without further VPATH interception. (See Figure 4-13.)

Figure 4-13 VPATH Interception

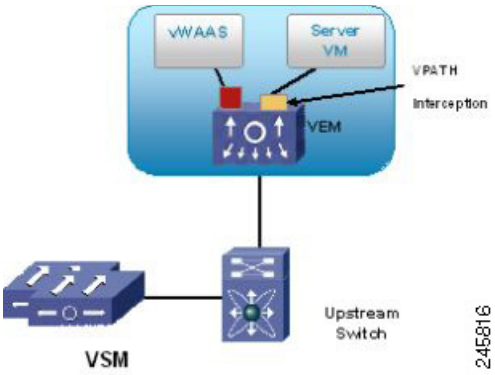

Interception is configured on the server VM port profile in both directions.

To configure VPATH interception on a vWAAS device, follow these steps:

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

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

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

The Device Dashboard window appears.

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

The VPATH settings window appears. (See Figure 4-14.)

Figure 4-14 VPATH Settings Window

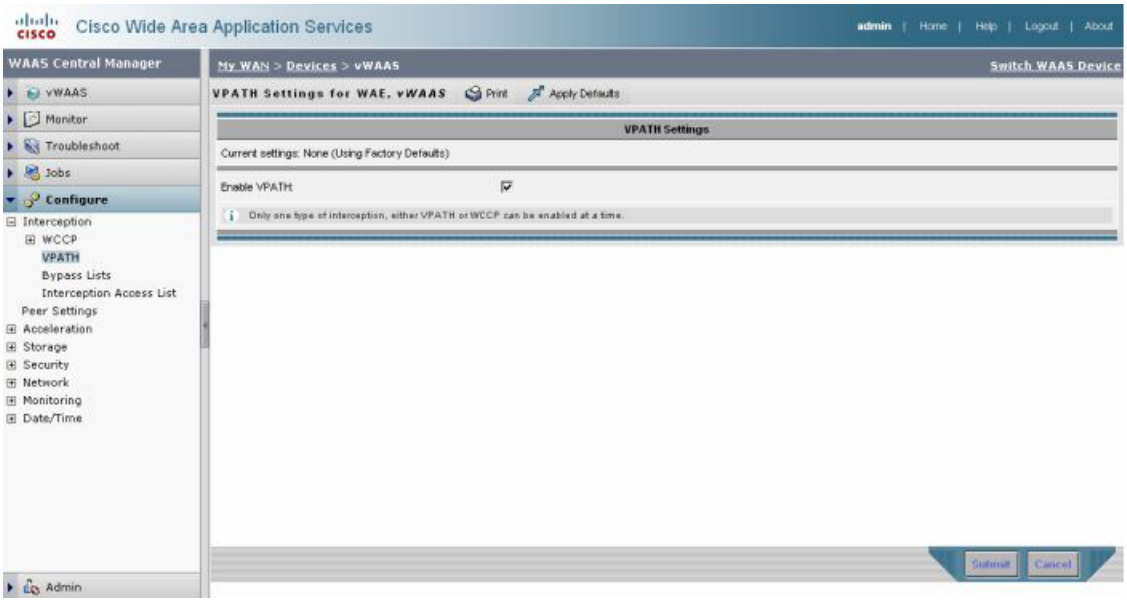

Check the **Enable VPATH** check box to enable VPATH interception on the vWAAS device.

Note Only one type of interception can be enabled at a time (VPATH or WCCP).

For device groups, VPATH interception is not applicable for vWAAS devices (OE-VWAAS) in the groups that have WCCP enabled and such devices go to override. VPATH settings also are not applicable to any physical WAAS devices in the group.

Step 4 Click **Submit**.

To enable VPATH from the CLI, use the **vn-service vpath** global configuration command. The default is disabled. For monitoring and troubleshooting, use the **show statistics vn-service vpath** and **clear statistics vn-service vpath** EXEC configuration commands.

For more information on virtual WAAS configuration, see the *Cisco Wide Area Application Services vWAAS Installation and Configuration Guide*.