CHAPTER 5

Configuring Network Settings

This chapter describes how to configure basic network settings such as creating additional network interfaces to support network traffic, specifying a DNS server, enabling Cisco Discovery Protocol (CDP), and configuring the directed mode of operation where peer WAEs exchange traffic using UDP encapsulation to avoid firewall traversal issues.

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

This chapter contains the following sections:

Configuring Network Interfaces

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

This section contains the following topics:

- Configuring Port-Channel Settings, page 5-8
- Configuring Interfaces for DHCP, page 5-9

We recommend that you use the WAAS Central Manager GUI instead of the WAAS CLI to configure network settings, but if you want to use the CLI, see the following commands in the *Cisco Wide Area Application Services Command Reference*: **interface**, **ip address**, **port-channel**, and **primary-interface**.

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

Configuring a Standby Interface

In this procedure, you configure a logical interface called a standby interface. After you set up the parameters for this logical interface, you must associate physical interfaces with the standby interface to create the standby group. (A standby group consists of two or more physical interfaces.) In the WAAS Central Manager GUI, you create the standby group by assigning a standby group priority to the physical interface. (See the ""Configuring the Interface Priority Setting" section on page 5-4.")

Standby interfaces remain inactive unless an active interface fails. When an active network interface fails (because of cable trouble, Layer 2 switch failure, high error count, or other failure), and that interface is part of a standby group, a standby interface can become active and take the load off the failed interface. With standby interface configuration, only one interface is active at a given time.

To configure standby interfaces, you must assign each physical interface to a standby group. The following rules define standby group relationships:

- A standby group consists of two or more physical interfaces.
- The maximum number of standby groups on a WAAS device is four.
- Each standby group is assigned a unique standby IP address, shared by all members of the group.
- Configuring the duplex and speed settings of the standby group member interfaces provides better reliability.
- IP ACLs can be configured on physical interfaces that are members of a standby group.
- Each interface in a standby group is assigned a priority. The operational interface with the highest priority in a standby group is the active interface. Only the active interface uses the group IP address.
- If the active interface fails, the operational interface in its standby group that is assigned the next highest priority becomes active.
- If all the members of a standby group fail, then one recovers, the WAAS software brings up the standby group on the operational interface.
- For a standby interface to show a status of active, the interface must be able to ping its IP default gateway.
- The priority of an interface in a standby group can be changed at runtime. The interface that has the highest priority after this change becomes the new active interface. (The default action is to preempt the currently active interface if an interface with higher priority exists.)
- The **errors** option, which is disabled by default, defines the maximum number of errors allowed on the active interface before the interface is shut down and before the standby is brought up.

To configure a standby interface, follow these steps:

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

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

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

Figure 5-1 Network Interfaces for Device Window

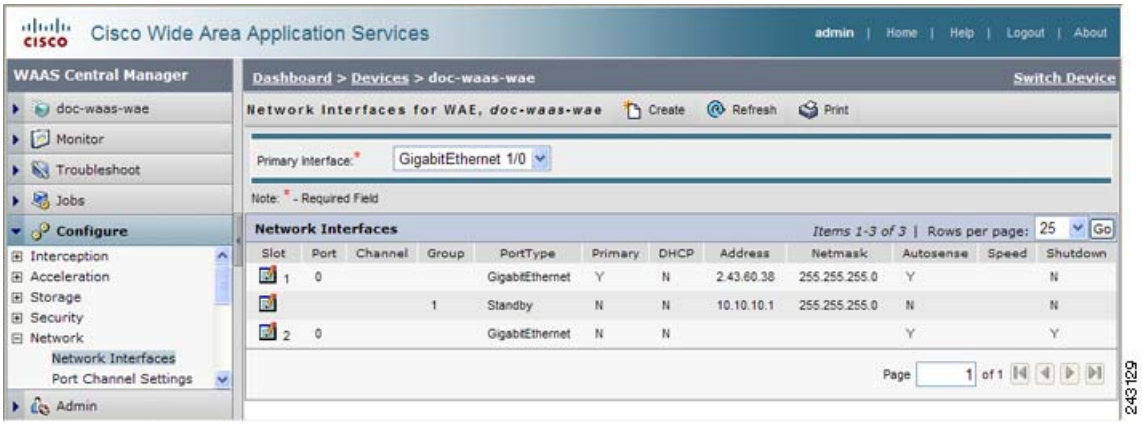

Step 4 In the taskbar, click the **Create New Interface** icon. The Creating New Network Interface window appears.

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

Step 6 From the Standby Group Number drop-down list, choose a group number (1–4) for the interface.

Step 7 In the Address field, specify the IP address of the standby group.

Step 8 In the Netmask field, specify the netmask of the standby group.

Step 9 In the Number of Errors field, enter the maximum number of errors allowed on this interface. The range is 0 to 4294967295.

Step 10 Check the **Shutdown** check box to shut down the hardware interface. By default, this option is disabled.

Step 11 In the Gateway field enter the default gateway IP address. If an interface is configured for DHCP, then this field is read only.

Step 12 Click **Submit**.

Step 13 Configure the interface priority setting as described in the ""Configuring the Interface Priority Setting" section on page 5-4."

Configuring the Interface Priority Setting

After you have configured a logical standby interface using the WAAS Central Manager GUI, you configure the standby group by setting a priority for each physical interface that you want to be associated with that standby group. The interface priority setting defines the active interface in a particular standby group and the order in which other interfaces in the standby group will become active if the active interface fails. The operational interface with the highest priority in a standby group is the active interface. Only the active interface uses the standby group IP address. You must have a standby interface configured before you can enter the priority settings in the WAAS Central Manager GUI. (See the """Configuring a Standby Interface" section on page 5-2.")

To configure the priority of the interface and associate it with a particular standby group, follow these steps:

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

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

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

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

Do not choose a logical interface (standby or portchannel) in this step. You cannot assign a standby priority to a logical interface.

Figure 5-2 Modifying Network Interface Window—Standby Group Priority Settings

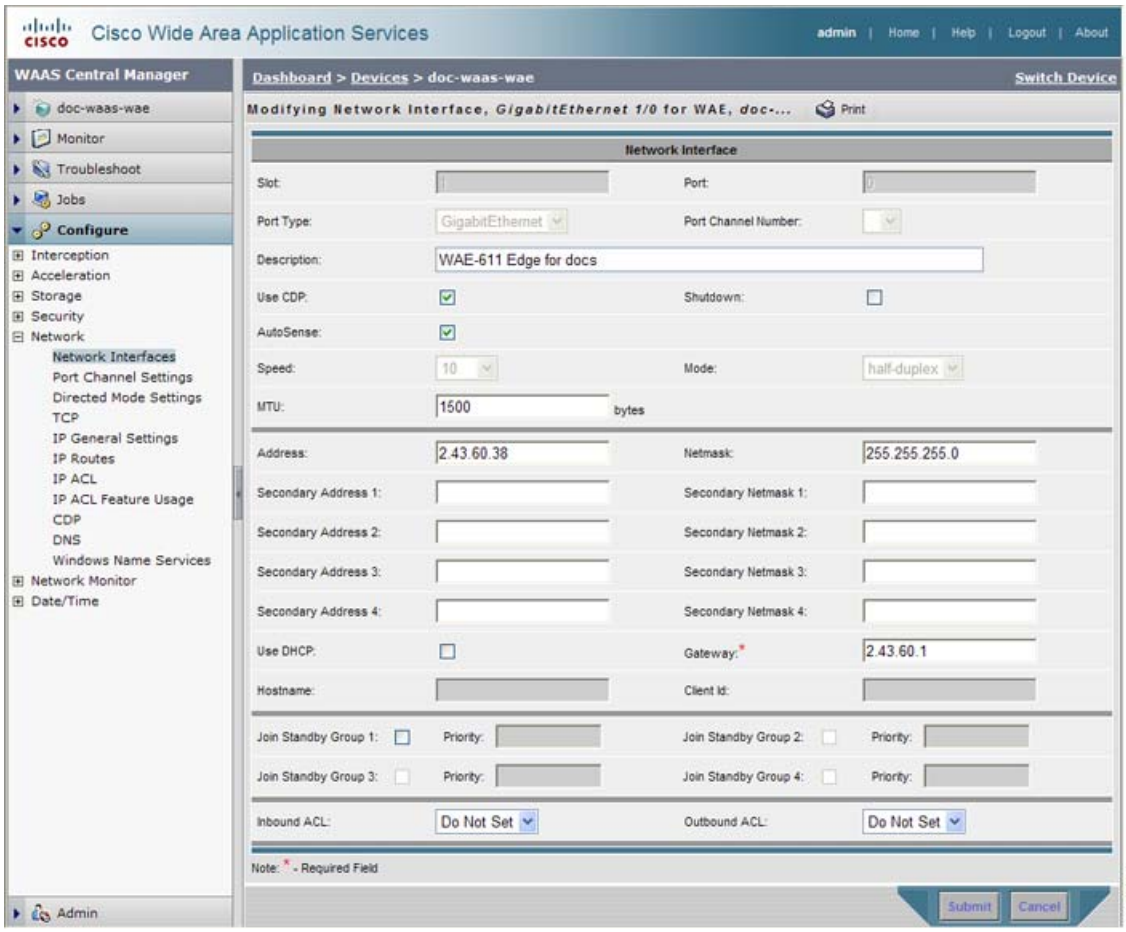

Step 5 Complete the following steps to specify the group and priority level number for this interface:

a. Scroll down the window until you see the Join Standby Group check boxes.

b. Check the **Join Standby Group** check box that you want this interface to join.

c. Enter a priority level number (0–4294967295) to set the priority of the interface in the standby group.

A Standby Group Priority field becomes available only when you have previously configured that standby group. (See the "Configuring a Standby Interface" section on page 5-2.) You can configure up to four standby groups for each WAAS device.

Note If an interface belongs to more than one standby group, you can configure the interface with a different priority in each standby group for better load balancing. For example, interfaces GE 0/0 and GE 0/1 are both in standby group 1 and in standby group 2. If you configure GE 0/0 with the highest priority in standby group 1 and configure GE 0/1 with the highest priority in standby group 2, standby group 1 will use GE 0/0 as the active interface, while standby group 2 will use GE 0/1 as the active interface. This configuration allows each interface to back up the other, if one of them fails.

Step 6 Click **Submit**. The interface joins the specified standby group.

Configuring Multiple IP Addresses on a Single Interface

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

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

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

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

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

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

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

Step 5 In the Secondary Address and Secondary Netmask fields 1 through 4, enter up to four different IP addresses and secondary netmasks for the interface.

Step 6 Click **Submit**.

Modifying Gigabit Ethernet Interface Settings

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

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

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

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

The Device Dashboard window appears.

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

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

Note On an NME-WAE device, the internal interface to the router is designated slot 1, port 0 and the external interface is designated slot 2, port 0. For NME-WAE configuration details, see the document *Configuring Cisco WAAS Network Modules for Cisco Access Routers*.

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

The Modifying Network Interface window appears, displaying the interface configurations on a particular slot and port.

Note Some of the fields in the window are not available. Interface configurations for slot, port, and port type are set for physical interfaces during initial startup or by using the WAAS CLI. The port channel number can be configured for a port channel interface when you create this type of interface in the WAAS Central Manager GUI; however, this field is not available when you modify a physical interface. (See the "Configuring Port-Channel Settings" section on page 5-8.)

Note When configuring the internal interface (GigabitEthernet 1/0) on an NME-WAE device, you cannot change the following fields or check boxes: AutoSense, Speed, Mode, Address, Netmask, and Use DHCP. If you attempt to change these values, the Central Manager displays an error when you click Submit. These settings for the internal interface can be configured only through the host router CLI. For details, see the document *Configuring Cisco WAAS Network Modules for Cisco Access Routers*.

Step 5 Check the **Use CDP** check box to enable Cisco Discovery Protocol (CDP) on an interface.

When enabled, CDP obtains protocol addresses of neighboring devices and discovers the platform of those devices. It also shows information about the interfaces used by your router.

Configuring CDP from the CDP Settings window enables CDP globally on all the interfaces. For information on configuring CDP settings, see the "Configuring CDP Settings" section on page 5-15.

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

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

Checking this check box disables the manual Speed and Mode drop-down list settings.

Note When autosense is on, manual configurations are overridden. You must reboot the WAAS device to start autosensing.

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

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

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

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

Full-duplex transmission allows data to travel in both directions at the same time through an interface or a cable. A half-duplex setting ensures that data only travels in one direction at any given time. Although full duplex is faster, the interfaces sometimes cannot operate effectively in this mode. If you encounter excessive collisions or network errors, you may configure the interface for half-duplex rather than full duplex.

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

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

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

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

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

Step 12 Click **Submit**.

Configuring Port-Channel Settings

WAAS software supports the grouping of up to four same-speed network interfaces into one virtual interface. This grouping capability allows you to set or remove a virtual interface that consists of two Gigabit Ethernet interfaces. This capability also provides interoperability with Cisco routers, switches, and other networking devices or hosts supporting EtherChannel, load balancing, and automatic failure detection and recovery based on each interface's current link status. EtherChannel is also referred to as a port channel.

To configure port-channel settings, follow these steps:

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

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

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

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

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

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

Step 6 In the Port Channel Number field, enter either **1** or **2** for the port-channel interface number.

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

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

Step 9 In the Address field, specify the IP address of the interface.

Step 10 In the Netmask field, specify the netmask of the interface.

Step 11 (Optional) From the Inbound ACL drop-down list, choose an IP ACL to apply to inbound packets.

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

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

Step 13 Click **Submit**.

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

Configuring Interfaces for DHCP

Note You must disable autoregistration before you can manually configure an interface for DHCP.

A WAAS device sends its configured client identifier and hostname to the DHCP server when requesting network information. You can configure DHCP servers to identify the client identifier information and the hostname information that the WAAS device is sending and then to send back the specific network settings that are assigned to the WAAS device.

To enable an interface for DHCP, follow these steps:

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

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

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

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

Note Do not choose a logical interface (standby or port channel) in this step, because you cannot configure DHCP on a logical interface. In addition, do not choose the internal interface (GigabitEthernet 1/0) on an NME-WAE device, because this interface can be configured only through the host router CLI. For details, see the document *Configuring Cisco WAAS Network Modules for Cisco Access Routers*.

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

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

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

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

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

Step 8 Click **Submit**.

Configuring a Load-Balancing Method for Interfaces

Before you configure load balancing, ensure that you have configured the port-channel settings described in the "Configuring Port-Channel Settings" section on page 5-8.

To configure load balancing, follow these steps:

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

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

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

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

- **dst-ip**—Destination IP address. .
- **round robin**—Each interface in the channel group. Round robin allows traffic to be distributed evenly among all interfaces in the channel group. The other balancing option gives you the flexibility to choose specific interfaces (by IP address) when sending an Ethernet frame. This option is selected by default.

Step 5 Click **Submit**.

To configure a load-balancing method from the CLI, you can use the **port-channel** global configuration command.

Configuring TCP Settings

For data transactions and queries between client and servers, the size of windows and buffers is important, so fine-tuning the TCP stack parameters becomes the key to maximizing cache performance.

Because of the complexities involved in TCP parameters, be careful in tuning these parameters. In nearly all environments, the default TCP settings are adequate. Fine tuning of TCP settings is for network administrators with adequate experience and full understanding of TCP operation details.

To configure TCP settings, follow these steps:

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

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

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

Figure 5-3 TCP Settings Window

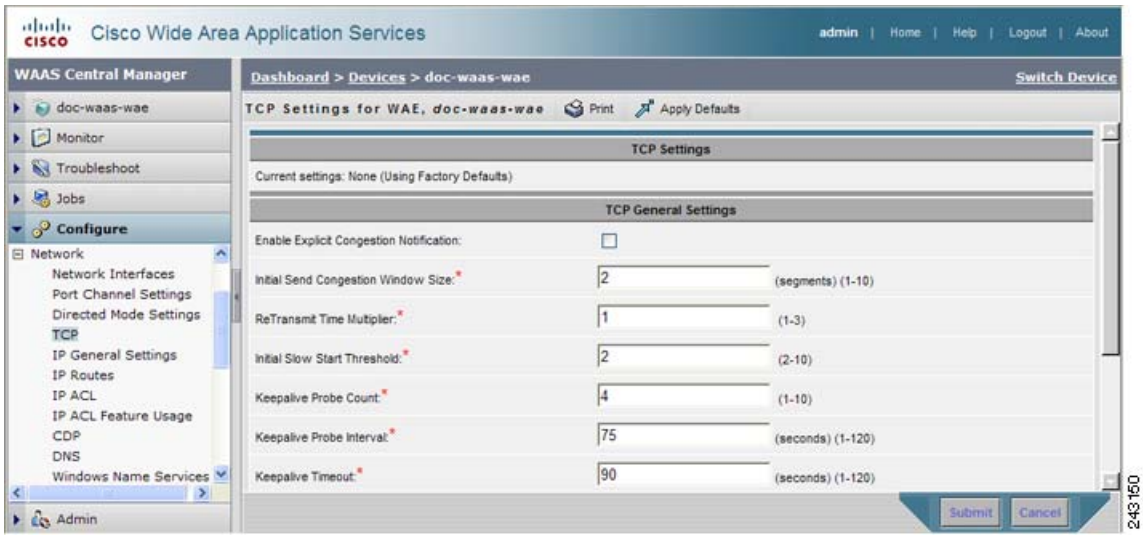

Step 4 Make the necessary changes to the TCP settings.

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

Step 5 Click **Submit**.

A "Click Submit to Save" message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking **Reset**. The Reset button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

Table 5-1 TCP Settings

TCP Setting	Description
TCP General Settings	
Enable Explicit Congestion Notification	Enables reduction of delay and packet loss in data transmissions. It provides TCP support for RFC 2581. By default, this option is disabled. For more information, see the "Explicit Congestion Notification" section on page 5-12.
Initial Send Congestion Window Size	Initial congestion window size value in segments. The range is 1 to 10 segments. The default is 2 segments. For more information, see the "Congestion Windows" section on page 5-12.

Table 5-1 TCP Settings (continued)

TCP Setting	Description
ReTransmit Time Multiplier	Factor used to modify the length of the retransmit timer by 1 to 3 times the base value determined by the TCP algorithm. The default is 1, which leaves the times unchanged. The range is 1 to 3. For more information, see the "Retransmit Time Multiplier" section on page 5-13. **Note** Modify this factor with caution. It can improve throughput when TCP is used over slow reliable connections but should never be changed in an unreliable packet delivery environment.
Initial Slow Start Threshold	Threshold for slow start in segments. The range is 2 to 10 segments. The default is 2 segments. For more information, see the "TCP Slow Start" section on page 5-13.
Keepalive Probe Count	Number of times that the WAAS device can retry a connection before the connection is considered unsuccessful. The range is 1 to 120 attempts. The default is 4 attempts.
Keepalive Probe Interval	Length of time that the WAAS device keeps an idle connection open. The default is 75 seconds.
Keepalive Timeout	Length of time that the WAAS device keeps a connection open before disconnecting. The range is 1 to 120 seconds. The default is 90 seconds.

To configure TCP settings from the CLI, you can use the **tcp** global configuration command.

This section contains the following topics:

Explicit Congestion Notification

The TCP Explicit Congestion Notification (ECN) feature allows an intermediate router to notify the end hosts of impending network congestion. It also provides enhanced support for TCP sessions associated with applications that are sensitive to delay or packet loss. The major issue with ECN is that the operation of both the routers and the TCP software stacks needs to be changed to accommodate the operation of ECN.

Congestion Windows

The congestion window (*cwnd*) is a TCP state variable that limits the amount of data that a TCP sender can transmit onto the network before receiving an acknowledgment (ACK) from the receiving side of the TCP transmission. The TCP *cwnd* variable is implemented by the TCP congestion avoidance algorithm. The goal of the congestion avoidance algorithm is to continually modify the sending rate so that the

sender automatically senses any increase or decrease in available network capacity during the entire data flow. When congestion occurs (manifested as packet loss), the sending rate is first lowered then gradually increased as the sender continues to probe the network for additional capacity.

Retransmit Time Multiplier

The TCP sender uses a timer to measure the time that has elapsed between sending a data segment and receiving the corresponding ACK from the receiving side of the TCP transmission. When this retransmit timer expires, the sender (according to the RFC standards for TCP congestion control) must reduce its sending rate. However, because the sender is not reducing its sending rate in response to network congestion, the sender is not able to make any valid assumptions about the current state of the network. Therefore, in order to avoid congesting the network with an inappropriately large burst of data, the sender implements the slow start algorithm, which reduces the sending rate to one segment per transmission. (See the ""TCP Slow Start" section on page 5-13.")

You can modify the sender's retransmit timer by using the Retransmit Time Multiplier field in the WAAS Central Manager GUI. The retransmit time multiplier modifies the length of the retransmit timer by one to three times the base value, as determined by the TCP algorithm that is being used for congestion control.

When making adjustments to the retransmit timer, be aware that they affect performance and efficiency. If the retransmit timer is triggered too early, the sender pushes duplicate data onto the network unnecessarily; if the timer is triggered too slowly, the sender remains idle for too long, unnecessarily slowing data flow.

TCP Slow Start

Slow start is one of four congestion control algorithms used by TCP. The slow start algorithm controls the amount of data being inserted into the network at the beginning of a TCP session when the capacity of the network is not known.

For example, if a TCP session began by inserting a large amount of data into the network, much of the initial burst of data would likely be lost. Instead, TCP initially transmits a modest amount of data that has a high probability of successful transmission. Next, TCP probes the network by sending increasing amounts of data as long as the network does not show signs of congestion.

The slow start algorithm begins by sending packets at a rate that is determined by the congestion window, or *cwnd* variable. (See the ""Congestion Windows" section on page 5-12.") The algorithm continues to increase the sending rate until it reaches the limit set by the slow start threshold (*ssthresh*) variable. Initially, the value of the *ssthresh* variable is adjusted to the receiver's maximum segment size (RMSS). However, when congestion occurs, the *ssthresh* variable is set to half the current value of the *cwnd* variable, marking the point of the onset of network congestion for future reference.

The starting value of the *cwnd* variable is set to that of the sender maximum segment size (SMSS), which is the size of the largest segment that the sender can transmit. The sender sends a single data segment, and because the congestion window is equal to the size of one segment, the congestion window is now full. The sender then waits for the corresponding ACK from the receiving side of the transmission. When the ACK is received, the sender increases its congestion window size by increasing the value of the *cwnd* variable by the value of one SMSS. Now the sender can transmit two segments before the congestion window is again full and the sender is once more required to wait for the corresponding ACKs for these segments. The slow start algorithm continues to increase the value of the *cwnd* variable and therefore

increase the size of the congestion window by one SMSS for every ACK received. If the value of the *cwnd* variable increases beyond the value of the *ssthresh* variable, then the TCP flow control algorithm changes from the slow start algorithm to the congestion avoidance algorithm.

Enabling the MTU Discovery Utility

The WAAS software supports the IP Path Maximum Transmission Unit (MTU) Discovery method, as defined in RFC 1191. When enabled, the Path MTU Discovery feature discovers the largest IP packet size allowable between the various links along the forwarding path and automatically sets the correct value for the packet size. By using the largest MTU that the links can handle, the sending device can minimize the number of packets it must send.

IP Path MTU Discovery is useful when a link in a network goes down, which forces the use of another, different MTU-sized link. IP Path MTU Discovery is also useful when a connection is first being established, and the sender has no information about the intervening links.

Note IP Path MTU Discovery is a process initiated by the sending device. If a server does not support IP Path MTU Discovery, the receiving device will have no available means to avoid fragmenting datagrams generated by the server.

By default, this feature is disabled. With the feature disabled, the sending device uses a packet size that is the lesser of 576 bytes and the next hop MTU. Existing connections are not affected when this feature is turned on or off.

To enable the MTU Discovery feature, follow these steps:

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

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

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

Step 4 Under the IP General Settings heading, enable the MTU discovery feature by checking the **Enable Path MTU Discovery** check box. By default, this option is disabled.

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

To enable the MTU discovery utility from the CLI, you can use the **ip path-mtu-discovery enable** global configuration command.

Configuring Static IP Routes

The WAAS software allows you to configure a static route for a network or host. Any IP packet designated for the specified destination uses the configured route.

To configure a static IP route, follow these steps:

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

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

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

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

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

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

Step 7 In the Gateway's IP Address field, enter the IP address of the gateway interface.

The gateway interface IP address should be in the same network as that of one of the device's network interfaces.

Step 8 Click **Submit**.

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

Aggregating IP Routes

An individual WAE device can have IP routes defined and can belong to device groups that have other IP routes defined.

In the IP Route Entries window, the Aggregate Settings radio button controls how IP routes are aggregated for an individual device, as follows:

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

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

Configuring CDP Settings

Cisco Discovery Protocol (CDP) is a device discovery protocol that runs on all Cisco-manufactured devices. With CDP, each device in a network sends periodic messages to all other devices in the network. All devices listen to periodic messages that are sent by others to learn about neighboring devices and determine the status of their interfaces.

With CDP, network management applications can learn the device type and the Simple Network Management Protocol (SNMP) agent address of neighboring devices. Applications are able to send SNMP queries within the network. CiscoWorks2000 also discovers the WAAS devices by using the CDP packets that are sent by the WAAS device after booting.

To perform device-related tasks, the WAAS device platform must support CDP to be able to notify the system manager of the existence, type, and version of the WAAS device platform.

To configure CDP settings, follow these steps:

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

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

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

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

Step 5 In the Hold Time field, enter the time (in seconds) to specify the length of time that a receiver is to keep the CDP packets.

The range is 10 to 255 seconds. The default is 180 seconds.

Step 6 In the Packet Send Rate field, enter a value (in seconds) for the interval between CDP advertisements.

The range is 5 to 254 seconds. The default is 60 seconds.

Step 7 Click **Submit**.

To configure CDP settings from the CLI, you can use the **cdp** global configuration command.

Configuring the DNS Server

DNS allows the network to translate domain names entered in requests into their associated IP addresses. To configure DNS on a WAAS device, you must complete the following tasks:

- Specify the list of DNS servers, which are used by the network to translate requested domain names into IP addresses that the WAAS device should use for domain name resolution.
- Enable DNS on the WAAS device.

To configure DNS server settings for a WAAS device, follow these steps:

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

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

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

Step 4 In the Local Domain Name field, enter the name of the local domain. You can configure up to three local domain names. Separate items in the list with a space.

Step 5 In the List of DNS Servers field, enter a list of DNS servers used by the network to resolve hostnames to IP addresses.

You can configure up to three DNS servers. Separate items in the list with a space.

Step 6 Click **Submit**.

A "Click Submit to Save" message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default and device group settings. To revert to the previously configured window settings, click **Reset**. The Reset button appears only when you have applied default or group settings to change the current device settings but the settings have not yet been submitted.

To configure DNS name servers from the CLI, you can use the **ip name-server** global configuration command.

Configuring Windows Name Services

To configure Windows name services for a device or device group, follow these steps:

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

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

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

Step 4 In the Workgroup or Domain Name field, enter the name of the workgroup (or domain) in which the chosen device or device group resides.

This name must be entered in shortname format and cannot exceed 127 characters. Valid characters include alphanumeric characters, a forward slash (\), an underscore (_), and a dash (-).

For example, if your domain name is cisco.com, the short name format is cisco.

Step 5 Check the **NT** check box if the workgroup or domain is a Windows NT 4 domain. For example, if your domain name is cisco.com, the short name format is cisco. If your workgroup or domain is a Windows 2000 or Windows 2003 domain, do not check the NT check box. By default, this option is disabled.

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

Step 7 Click **Submit**.

To configure Windows name services from the CLI, you can use the **windows-domain** global configuration command.

Configuring Directed Mode

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

Any firewall between two WAE peers must be configured to pass UDP traffic on port 4050, or whatever custom port is configured for directed mode if a port other than the default is used. Additionally, because the WAAS automatic discovery process uses TCP options before directed mode begins sending UDP traffic, the firewall must be configured to pass the TCP options. Cisco firewalls can be configured to allow TCP options by using the **ip inspect waas** command (for IOS 12.4(11)T2 and later) or the **inspect waas** command (for FWSM 3.2(1) and later and PIX 7.2(3) and later).

After directed mode is activated, the WAE transparently intercepts only packets coming from the LAN, while WAN packets are directly routed between the WAEs using UDP.

Directed mode operates with all configurable methods of traffic interception. Directed mode requires that you configure the WAAS devices (or Cisco WAE Inline Network Adapters) with routable, non-NATed IP addresses. When using directed mode with inline mode, you must configure the Cisco WAE Inline Network Adapter with routable IP addresses on its interfaces or traffic is black holed.

If a WAE at either end of a peer WAE connection specifies directed mode, and both WAEs support directed mode, then both WAEs use directed mode, even if one is not explicitly configured for directed mode. If a peer WAE does not support directed mode, then the peers pass through traffic unoptimized and each WAE creates a transaction log entry that notes the failed directed mode attempt.

You can invoke directed mode operation in the following ways:

- Directed mode can be explicitly activated in the WAAS Central Manager GUI or CLI.
- Directed mode can be automatically invoked when a peer WAE requests that directed mode be used.

To activate directed mode, follow these steps:

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

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

Step 3 In the navigation pane, choose **Configure > Network > Directed Mode Settings**. The Directed Mode Settings window appears. (See Figure 5-4.)

Figure 5-4 Directed Mode Settings Window

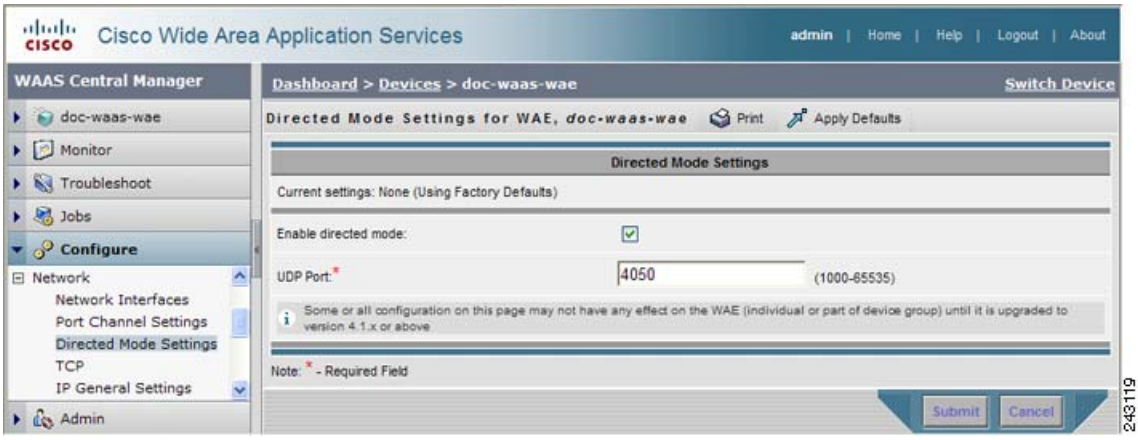

Step 4 Check the **Enable** check box to activate directed mode.

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

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

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