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, and enabling Cisco Discovery Protocol (CDP).

Note Throughout this chapter, the term WAAS device is used to refer collectively to the WAAS Central Managers and WAEs in your network.

This chapter contains the following topics:

- Configuring Network Interfaces, page 5-1
- Configuring a Load-Balancing Method for Interfaces, page 5-10
- Configuring TCP Settings, page 5-10
- Enabling the MTU Discovery Utility, page 5-14
- Configuring Static IP Routes, page 5-15
- Configuring CDP Settings, page 5-16
- Configuring the DNS Server, page 5-17
- Configuring Windows Name Services, page 5-18

Configuring Network Interfaces

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

This section contains the following topics:

- Configuring a Standby Interface, page 5-2
- Configuring the Interface Priority Setting, page 5-3
- Configuring Multiple IP Addresses on a Single Interface, page 5-5
- Modifying Gigabit Ethernet Interface Settings, page 5-5
- Modifying the Fibre Channel Interface, page 5-7
- 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**.

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 "Configuring the Interface Priority Setting.")

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 interface is assigned a unique IP address, and 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.
- 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.

Note Interface IP addresses and standby group IP addresses must be on different subnets to ensure reliable operation. You can use dummy IP addresses in the private address space to serve as interface primary IP addresses and use the real WAAS device IP address to serve as the standby group IP address in a different subnet to satisfy this requirement.

To configure a standby interface, follow these steps:

Step 1 From the WAAS Central Manager GUI, choose **Devices > 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 Home window appears.

Step 3 Click **Show Advanced** to display all the menu items in the Contents pane.

Step 4 In the Contents pane, choose **General Settings > Network > Network Interfaces**. The Network Interfaces window for the device appears.

Step 5 In the taskbar, click the **Create New Interface** icon. The Creating New Network Interface window appears.

Step 6 From the Port Type drop-down list, choose **Standby**. The window refreshes with fields for configuring the standby group settings.

Step 7 From the Standby Group Number drop-down list, choose a group number (1–4) for the interface.

Step 8 In the Address field, specify the IP address of the standby group.

Step 9 In the Netmask field, specify the netmask of the standby group.

Step 10 In the Number of Errors field, enter the maximum number of errors allowed on this interface. The range is 0 to 4294967295.

Step 11 Check the **Shutdown** check box to shut down the hardware interface. By default, this option is disabled.

Step 12 In the Gateway field enter the default gateway IP address. If an interface is configured for DHCP, then this field is read only.

Step 13 Click **Submit**.

Step 14 Configure the interface priority setting as described in "Configuring the Interface Priority Setting."

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.")

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

Step 1 From the WAAS Central Manger GUI, choose **Devices > 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 Home window appears.

Step 3 Click **Show Advanced** to display all the menu items in the Contents pane.

Step 4 In the Contents pane, choose **General Settings > Network > Network Interfaces**. The Network Interfaces window for the device appears.

Step 5 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 5-1.)

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

Figure 5-1 Modifying Network Interface Window—Standby Group Priority Settings

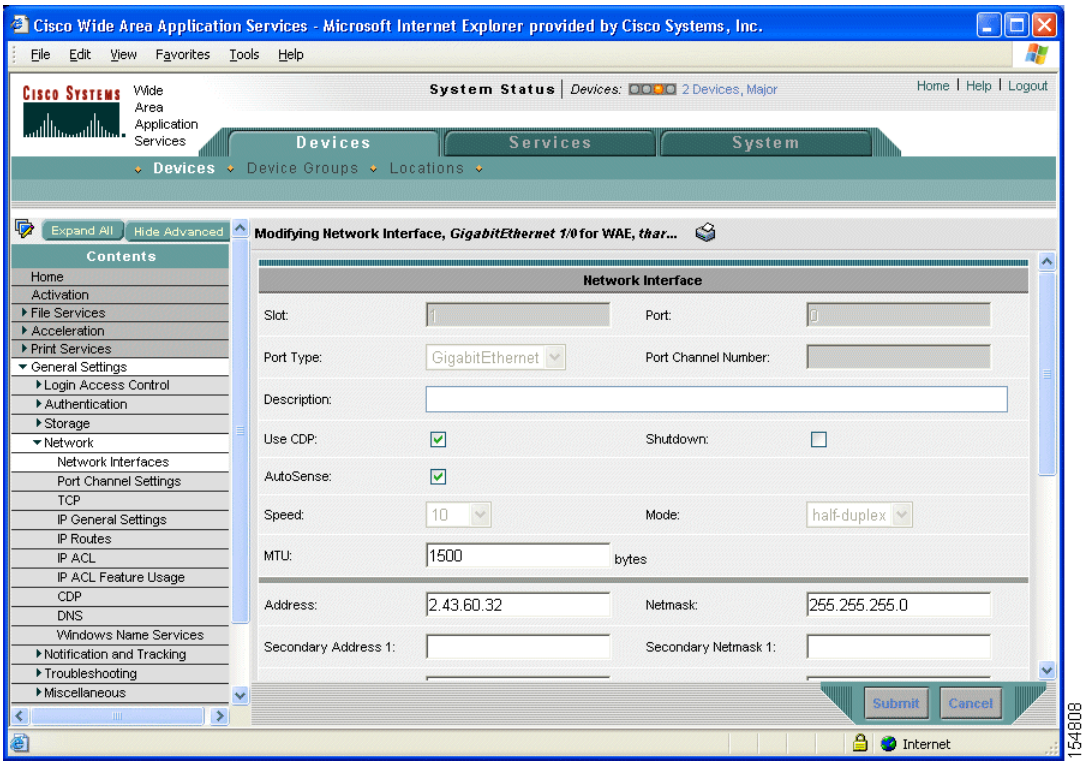

Step 6 Complete the following steps to specify the group and priority level number for this interface:

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 only becomes available 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 7 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, choose **Devices > 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 Home window appears.

Step 3 Click **Show Advanced** to display all the menu items in the Contents pane.

Step 4 In the Contents pane, choose **General Settings > Network > Network Interfaces**. The Network Interfaces listing window appears.

Step 5 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 portchannel) in this step. You cannot configure multiple interfaces on a logical interface.

Step 6 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 7 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, choose **Devices > 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 Home window appears with the Contents pane on the left.

Step 3 Click **Show Advanced** to display all the menu items in the Contents pane.

Step 4 In the Contents pane, choose **General Settings > Network > Network Interfaces**.

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

Step 5 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.)

Step 6 To enable Cisco Discovery Protocol (CDP) on an interface, check the **Use CDP** check box.

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 5-16.

Step 7 Check the **Shutdown** check box to shut down the hardware interface.

Step 8 To set the interface to autonegotiate the speed and mode, check the **AutoSense** check box.

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

Note Cisco router Ethernet interfaces do not negotiate duplex settings. If the WAAS device is connected to a router directly with a crossover cable, the device interface must be manually set to match the router interface settings. Disable autosense before configuring an Ethernet interface. When autosense is on, manual configurations are overridden. You must reboot the WAAS device to start autosensing.

Step 9 To manually configure the interface transmission speed and mode settings, follow these steps:

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, configure the interface for half-duplex rather than full duplex.

Step 10 Specify a value (in bytes) in the MTU field to set the interface Maximum Transmission Unit (MTU) size.

The range is 68–1500 bytes. The MTU is the largest size of IP datagram that can be transferred using a specific data link connection.

Step 11 Enter a new IP address in the Address field to change the interface IP address.

Step 12 Enter a new netmask in the Netmask field to change the interface netmask.

Step 13 Click **Submit**.

Modifying the Fibre Channel Interface

WAAS software supports Fibre Channel interfaces. Fibre Channel is the chosen technology for interconnecting storage devices and servers in a storage area network (SAN). In a SAN, the storage need not be directly attached to the server, and data transfer occurs over a high-throughput, high-availability network. Fibre Channel is capable of operating at speeds of 1 gigabit per second (Gbps) and 2 Gbps.

To detect the presence of Fibre Channel storage, the Fibre Channel array must be configured to assign storage space for the WAAS device, and the WAAS device must be reloaded before it can detect the storage assignment. To confirm whether or not the WAAS device has detected the storage assignment, use the **show disks** and the **show disks details** commands.

Note Logical interfaces, such as standby and Port Channel interfaces, can be created, removed, and modified by using the WAAS Central Manager GUI. Physical interfaces, such as Gigabit Ethernet and Fibre Channel interfaces, can only be modified.

To modify the settings of an existing Fibre Channel interface for the WAAS device, follow these steps:

Step 1 From the WAAS Central Manager GUI, choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the WAAS network.

Step 2 Click the **Edit** icon next to the WAAS device for which you want to modify the Fibre Channel interface settings. The Device Home window appears with the Contents pane on the left.

Step 3 Click **Show Advanced** to display all the menu items in the Contents pane.

Step 4 In the Contents pane, choose **General Settings > Network > Network Interfaces**. The Network Interfaces window appears, listing the network interfaces configured on particular ports and slots.

Step 5 Click the **Edit Network Interface** icon next to the Fibre Channel interface that you want to modify.

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

The following fields in the window are *not* available because they were configured using the **interface** configuration command:

- Slot—Slot number for the selected interface.
- Port—Port number for the selected interface.
- Port Type—Type of Fibre Channel port configured on the WAAS device.

Step 6 From the Speed drop-down list, choose one of the following speed options for the specified interface:

- autosense—Sets the Fibre Channel interface to automatically sense the interface speed.
- 1—Sets the Fibre Channel interface speed to 1 Gbps.
- 2—Sets the Fibre Channel interface speed to 2 Gbps.

Step 7 From the Mode drop-down list, choose one of the following modes options:

- autosense—Sets the operation mode of the WAAS device to autosense.
- direct-attached—Sets the operation mode when the WAAS device is directly connected to a storage array.
- switched—Sets the operation mode when the WAAS device is connected to a switch.

Step 8 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 the setting or removing of 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, choose **Devices > 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 Home window appears.

Step 3 Click **Show Advanced** to display all the menu items in the Contents pane.

Step 4 In the Contents Pane, choose **General Settings > Network > Network Interfaces**. The Network Interfaces window appears, listing all the interfaces for the chosen device.

Step 5 In the taskbar, click the **Create New Interface** icon. The Creating New Network Interface window appears.

Step 6 From the Port Type drop-down list, choose **Port Channel**.

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

Step 7 In the Port Channel Number field, enter either **1** or **2** for the port channel interface number.

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

Step 9 In the Gateway field, enter the default gateway IP address.

Step 10 In the Address field, specify the IP address of the interface.

Step 11 In the Netmask field, specify the netmask of the interface.

Step 12 (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 13 (Optional) From the Outbound ACL drop-down list, choose an IP ACL to apply to outbound packets.

Step 14 Click **Submit**.

Configuring Interfaces for DHCP

Note Autoregistration must be disabled before you can manually configure an interface for DHCP.

A WAAS device sends its configured client identifier and host name to the DHCP server when requesting network information. DHCP servers can be configured to identify the client identifier information and the host name information that the WAAS device is sending, then 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, choose **Devices > 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 Home window appears.

Step 3 Click **Show Advanced** to display all the menu items in the Contents pane.

Step 4 In the Contents pane, choose **General Settings > Network > Network Interfaces**. The Network Interfaces listing window appears.

Step 5 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 portchannel) in this step. You cannot configure DHCP on a logical interface.

Step 6 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 7 In the Hostname field, specify the hostname for the WAAS device or other device.

Step 8 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 9 Click **Submit**.

Configuring a Load-Balancing Method for Interfaces

Before you configure load-balancing, make sure 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, choose **Devices > Devices** or **Devices > 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 Click **Show Advanced** to display all the menu items in the Contents pane.

Step 4 In the Contents Pane, choose **General Settings > Network > Port Channel Settings**.

Step 5 From the Load Balancing Method drop-down list, choose a load-balancing method:

- dst-ip—Destination IP address.
- dst-mac—Destination MAC 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 options give you the flexibility to choose specific interfaces (by IP address or MAC address) when sending an Ethernet frame. This option is selected by default.

Step 6 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. The TCP memory limit settings allow you to control the amount of memory that can be used by the TCP subsystem send and receive buffers.

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.

Caution Do not modify the default TCP memory limit values unless you are knowledgeable about the changes you want to make. The default values are device dependent and have been chosen after extensive testing. They should not be changed under normal conditions. Increasing these values can result in the TCP subsystem using more memory, which might cause the system to be unresponsive. Decreasing these values can result in increased response times and lower performance.

To configure TCP settings, follow these steps:

Step 1 From the WAAS Central Manager GUI, choose **Devices > Devices** or **Devices > 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 Home window appears.

Step 3 Click **Show Advanced** to display all the menu items in the Contents pane.

Step 4 In the Contents pane, choose **General Settings > Network > TCP**. The TCP Settings window appears. (See Figure 5-2.)

Figure 5-2 TCP Settings Window

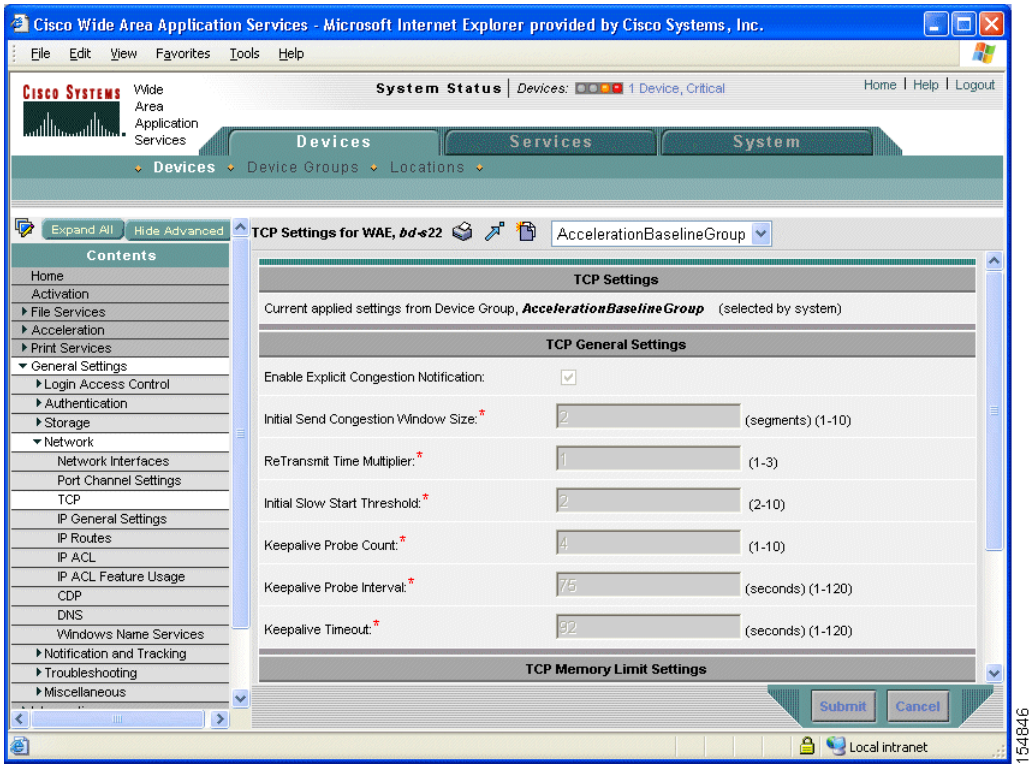

Step 5 Make the necessary changes to the TCP settings.

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

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 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 "About Explicit Congestion Notification" section on page 5-13.
Initial Send Congestion Window Size	Initial congestion window size value in segments. The default is 2 segments. For more information, see the "About Congestion Windows" section on page 5-13.
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. For more information, see the "About 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 default is 2 segments. For more information, see the "About TCP Slow Start" section on page 5-14.
Keepalive Probe Count	Number of times the WAAS device can retry a connection before the connection is considered unsuccessful. 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 default is 90 seconds.
TCP Memory Limit Settings	
TCP Limit Low Water Mark	The lower limit (in MB) of memory pressure mode, below which TCP enters into normal memory allocation mode. The range is 4 to 600. The low water mark must be a number that is less than the high water mark pressure setting.
TCP Memory Limit High Water Mark–Pressure	The upper limit (in MB) of normal memory allocation mode, beyond which TCP enters into memory pressure mode. The range is 5 to 610. The high water mark pressure must be a number that is less than the high water mark absolute setting
TCP Memory Limit High Water Mark–Absolute	The absolute limit (in MB) on TCP memory usage. The range is 6 to 620.

Table 5-2 describes the default values for each TCP memory limit setting, which are based on the total amount of memory for the device.

Table 5-2 Default TCP Memory Limit Settings

Total System Memory	Low	Pressure	Absolute
1 GB, 2 GB, or 4 GB	360 MB	380 MB	400 MB
512 MB	180 MB	190 MB	200 MB
256 MB	25 MB	28 MB	30 MB

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

About 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.

About 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.

About the 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 "About TCP Slow Start.")

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.

About 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 "About Congestion Windows.") 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

Cisco 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, choose **Devices > Devices** or **Devices > Device Groups**.

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

Step 3 Click **Show Advanced** to display all the menu items in the Contents pane.

Step 4 In the Contents pane, choose **General Settings > Network > IP General Settings**. The IP General Settings window appears.

Step 5 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 6 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

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, choose **Devices > Devices** or **Devices > Device Groups**.

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

Step 3 Click **Show Advanced** to display all the menu items in the Contents pane.

Step 4 In the Contents pane, choose **General Settings > Network > IP Routes**. The IP Route Entries window appears.

Step 5 In the taskbar, click the **Create New IP Route Entry** icon. The Creating New IP Route window appears.

Step 6 In the Destination Network Address field, enter the destination network IP address.

Step 7 In the Netmask field, enter the destination host netmask.

Step 8 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 9 Click **Submit**.

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

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, choose **Devices > Devices** or **Devices > Device Groups**.

Step 2 Click the **Edit** icon next to the device or device group that you want to configure. The Contents pane appears on the left.

Step 3 Click **Show Advanced** to display all menu items in the Contents pane.

Step 4 From the Contents pane, choose **General Settings > Network > CDP**. The CDP Settings window appears. (See Figure 5-3.)

Figure 5-3 CDP Settings Window

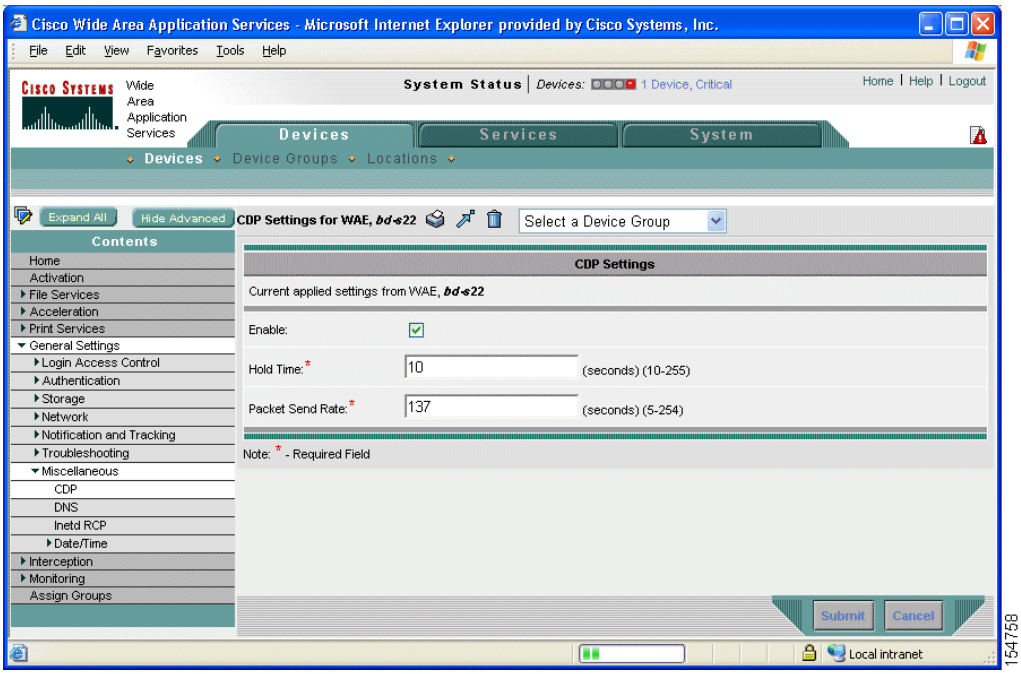

Step 5 Check the **Enable** check box to enable CDP support. By default, this option is enabled.

Step 6 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 7 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 8 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, choose **Devices > Devices** or **Devices > Device Groups**.

Step 2 Click the **Edit** icon next to the device or device group that you want to configure. The Contents pane appears on the left.

Step 3 From the Contents pane, choose **General Settings > 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 host names 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, choose **Devices > Devices** or **Devices > 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 Click **Show Advanced** to display all the menu items in the Contents pane.

Step 4 From the Contents pane, choose **General Settings > Network > Windows Name Services**. The Windows Name Services Settings window appears.

Step 5 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 6 Check the **NT Domain** 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 Domain check box. By default, this option is disabled.

Step 7 In the WINS server field, enter the hostname or IP address of the Windows Internet Naming Service (WINS) server.

Step 8 Click **Submit**.

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