Configuring the Gateway to Receive Email

This chapter contains the following sections:

- Overview of Configuring the Gateway to Receive Email, on page 1
- Working with Listeners, on page 2
- Configuring Global Settings for Listeners, on page 4
- Listening for Connection Requests by Creating a Listener Using Web Interface, on page 7
- Listening for Connection Requests by Creating a Listener Using CLI, on page 12
- Enterprise Gateway Configuration, on page 14

Overview of Configuring the Gateway to Receive Email

It is recommended that you avoid adding, changing, or deleting listeners on Cisco Secure Email Gateways.

The email gateway functions as the gateway for your organization, servicing email connections, accepting messages, and relaying them to the appropriate systems. The email gateway can service email connections from the Internet to recipients hosts inside your network, and from systems inside your network to the Internet. Typically, email connection requests use Simple Mail Transfer Protocol (SMTP). The appliance services SMTP connections by default, and acts as the SMTP gateway, also known as a mail exchanger or "MX," for the network.

The email gateway uses *listeners* to service incoming SMTP connection requests, A listener describes an email processing service that is configured on a particular IP interface. Listeners apply to email entering the appliance, from either the Internet or from systems within your network trying to reach the Internet. Use listeners to specify criteria that messages and connections must meet in order to be accepted and for messages to be relayed to recipient hosts. You can think of a listener as an "SMTP daemon" running on a specific port for each IP address specified. Also, listeners define how the email gateway communicates with systems that try to send email to the email gateway.

You can create the following types of listeners:

- **Public.** Listens for and accepts email messages coming in from the Internet. Public listeners receive connections from many hosts and direct messages to a limited number of recipients.
- **Private.** Listens for and accepts email messages coming from systems within the network, typically from internal groupware and email servers (POP/IMAP), intended for recipients outside the network in the Internet. Private listeners receive connections from a limited (known) number of hosts and direct messages to many recipients.

When you create a listener, you also must specify the following information:

- **Listener properties.** Define global properties that apply to all listeners, and properties specific to each listener. For example, you can specify the IP interface and port to use for a listener, and whether it is a public or private listener. For details on how to do this, see Working with Listeners, on page 2.
- **Which hosts that are allowed to connect to the listener.** Define a set of rules that control incoming connections from remote hosts. For example, you can define remote hosts and whether or not they can connect to the listener. For details on how to do this, see Defining Which Hosts Are Allowed to Connect Using the Host Access Table.
- **(Public listeners only) The local domains for which the listener accepts messages.** Define which recipients are accepted by the public listener. For example, if your organization uses the domain `currentcompany.com` and it previously used `oldcompany.com` , then you might accept messages for both currentcompany.com and `oldcompany.com` . For details on how to do this, see Accepting or Rejecting Connections Based on Domain Name or Recipient Address.

The settings configured in the listener, including its Host Access Table and Recipient Access Table, affect how the listener communicates with an SMTP server during the SMTP conversation. This allows the email gateway to block a spamming host before the connection is closed.

Figure 1: Relationship between Listeners, IP Interfaces, and Physical Ethernet Interfaces

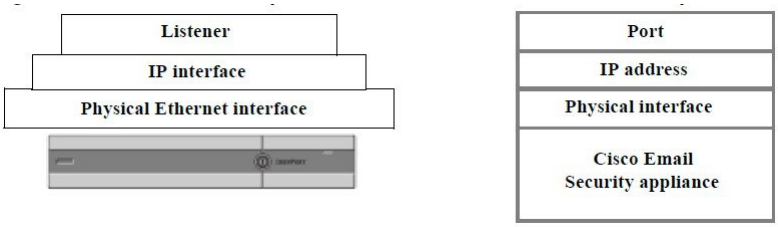

Working with Listeners

Configure listeners on the Network > Listeners page in the GUI, or using the listenerconfig command in the CLI.

You can define global settings that apply to all listeners. For more information, see Configuring Global Settings for Listeners, on page 4.

Consider the following rules and guidelines when working with and configuring listeners on the email gateway:

- You can define multiple listeners per configured IP interface, but each listener must use a different port.
- By default, listeners use SMTP as the mail protocol to service email connections. However, you can also configure the appliance to service email connections using Quick Mail Queuing Protocol (QMQP). Do this using the `listenerconfig` CLI command.
- Listeners support both Internet Protocol version 4 (IPv4) and version 6 (IPv6) addresses. You can use either protocol version or both on a single listener. The listener uses the same protocol version for mail delivery as the connecting host. For example, if the listener is configured for both IPv4 and IPv6 and connects to a host that uses IPv6, the listener uses IPv6. However, if the listener is configured to only use IPv6 addresses, it cannot connect to a host that is only using IPv4 addresses.
- At least one listener (with default values) is configured on the email gateway after running the System Setup Wizard. However, when you create a listener manually, AsyncOS does not use these default IP Reputation score values.

- **C170 and C190 appliances**: By default, the System Setup Wizard walks you through configuring one public listener for both receiving mail from the Internet and for relaying email from your internal network. That is, one listener can perform both functions.
- To help test and troubleshoot the email gateway, you can create a "sinkhole" type listener instead of a public or private listener. When you create a sinkhole listener, you choose whether messages are written to disk or not before they are deleted. (See the "Testing and Troubleshooting" chapter for more information.) Writing messages to disk before deleting them can help you measure the rate of receiving and the speed of the queue. A listener that doesn't write messages to disk can help you measure the pure rate of receiving from your message generation systems. This listener type is only available through the `listenerconfig` command in the CLI.

Figure - Public and Private Listeners on Email Gateway Models with More than Two Ethernet Interfaces illustrates a typical email gateway configuration created by the System Setup Wizard on email gateway models that have more than two Ethernet interfaces. Two listeners are created: a public listener to serve inbound connections on one interface and a private listener to serve outbound connections on a second IP interface.

Figure - Public Listener on Email Gateway Models with Only Two Ethernet Interfaces illustrates a typical email gateway configuration created by the System Setup Wizard on email gateway models that have only two Ethernet interfaces. One public listener on a single IP interface is created to serve both inbound and outbound connections.

Figure 2: Public and Private Listeners on Email Gateway Models with More than Two Ethernet Interfaces

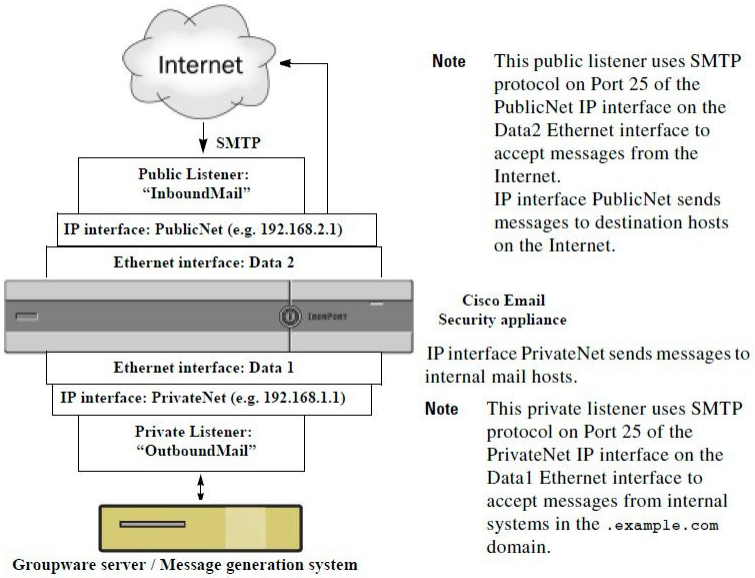

Figure 3: Public Listener on Email Gateway Models with Only Two Ethernet Interfaces

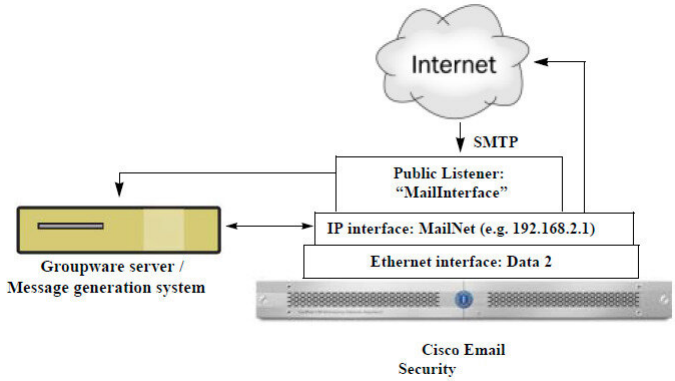

Note This public listener uses SMTP protocol on Port 25 of the PublicNet IP interface on the Data2 Ethernet interface to accept messages from the Internet and to relay messages from internal systems in the `.example.com` domain. IP interface MailNet sends messages to destination hosts on the Internet and to internal mail hosts

Configuring Global Settings for Listeners

Global settings for the listeners affect all of the listeners that are configured on the email gateway. If the listener uses an interface that has both Internet Protocol version 4 (IPv4) and version 6 (IPv6) addresses, the listener settings apply to both IPv4 and IPv6 traffic

Procedure

Step 1 Choose **Network > Listeners**.

Step 2 Click **Edit Global Settings**.

Step 3 Make changes to the settings defined in the following table.

Table 1: Listener Global Settings

Global Setting	Description
Maximum Concurrent Connections	Set the maximum number of concurrent connections for listeners. The default value is 300 for the C3x0 and C6x0 models, and the default value is 50 for the C1x0 models.. If the listener accepts both IPv4 and IPv6 connections, the number of connections is divided between the two. For example, if the maximum concurrent connections is 300, then the sum of IPv4 and IPv6 connections cannot exceed 300.
Maximum Concurrent TLS Connections	Set the maximum concurrent TLS connections across all listeners combined. The default value is 100. If the listener accepts both IPv4 and IPv6 TLS connections, the number of connections is divided between the two. For example, if the maximum concurrent connections is 100, then the sum of IPv4 and IPv6 TLS connections cannot exceed 100.

Global Setting	Description
Injection Counters Reset Period	Allows you to adjust when the injection control counters are reset. For very busy systems maintaining counters for a very large number of different IP addresses, configuring the counters to be reset more frequently (for example, every 15 minutes instead of every 60 minutes) will ensure that the data does not grow to an unmanageable size and impact system performance. The current default value is 1 hour. You can specify periods ranging from as little as 1 minute (60 seconds) to as long as 4 hours (14,400 seconds). See Injection Control Periodicity.
Timeout Period for Unsuccessful Inbound Connections	Set the length of time AsyncOS will allow an unsuccessful inbound connection to remain intact before closing it. An unsuccessful connection can be an SMTP conversation in which SMTP or ESMTP commands continue to be issued without a successful message injection occurring. When the specified timeout is reached, the behavior is to send an error and disconnect: "421 Timed out waiting for successful message injection, disconnecting." A connection is considered unsuccessful until it successfully injects a message. Only available for SMTP connections on public listeners. The default value is 5 minutes.
Total Time Limit for All Inbound Connections	Set the length of time AsyncOS will allow an inbound connection to remain intact before closing it. This setting is intended to preserve system resources by enforcing a maximum allowable connection time. Once about 80% of this maximum connection time is reached the following message is issued: "421 Exceeded allowable connection time, disconnecting." The email gateway will attempt to disconnect when the connection exceeds 80% of the maximum connection time in order to prevent disconnecting mid-message. It is likely that a problem is occurring with the inbound connection if it is open long enough to reach 80% of the maximum connection time. Keep this threshold in mind when specifying the time limit. Only available for SMTP connections on public listeners. The default value is 15 minutes.
Maximum size of subject	Messages having subject size within the specified limit will be accepted and any other messages will be rejected. If you set this value to 0, no limit is applied.

Global Setting	Description
HAT delayed rejections	Configure whether to perform HAT rejection at the message recipient level.By default, HAT rejected connections will be closed with a banner message at the start of the SMTP conversation. When an email is rejected due to HAT "Reject" settings, AsyncOS can perform the rejection at the message recipient level (RCPT TO), rather than at the start of the SMTP conversation. Rejecting messages in this way delays the message rejection and bounces the message, allowing AsyncOS to retain more detailed information about the rejected messages. For example, you can see the mail from address and each recipient address of the message which is blocked. Delaying HAT rejections also makes it less likely that the sending MTA will perform multiple retries. When you enable HAT delayed rejection, the following behavior occurs: The MAIL FROM command is accepted, but no message object is created. All RCPT TO commands are rejected with text explaining that access to send e-mail is refused. If the sending MTA authenticates with SMTP AUTH, they are granted a RELAY policy and are allowed to deliver mail as normal. Only configurable from the CLI `listenerconfig --> setup command`.

Step 4 Submit and commit your changes.

What to do next

Related Topics

- Settings for Messages Containing Multiple Encodings, on page 6

Settings for Messages Containing Multiple Encodings

You can define the behavior of the email gateway while modifying the message encoding for the following parameters:

- Headers
- Untagged non-ASCII headers
- Mismatched footer or heading encoding

To configure this behavior, use the localeconfig command in CLI.

Note This behavior cannot be configured using web interface.

For a sample CLI transcript, see Disclaimer Stamping and Multiple Encodings.

Listening for Connection Requests by Creating a Listener Using Web Interface

Procedure

Step 1 Choose `Network > Listener`.

Step 2 Click **Add Listener**.

Step 3 Configure the settings defined in the following table.

Table 2: Listener Settings

Name	Unique nickname you supply for the listener, for future reference. The names you define for listeners are case-sensitive. AsyncOS will not allow you to create two identical listener names.
Type of Listener	Choose one of the following types of listeners: • **Public.** Public listeners contain default characteristics for receiving email from the Internet. • **Private.** Private listeners are intended to be used for private (internal) networks.
Interface	Choose a configured email gateway IP interface and TCP port on which to create the listener. Depending on the version of the IP address used by the interface, the listener accepts connections from IPv4 addresses, IPv6 addresses or from both versions. By default, SMTP uses port 25 and QMQP uses port 628.
Bounce Profile	Select a bounce profile (bounce profiles created via the `bounceconfig` command in the CLI are available in the list, see Creating a New Bounce Profile).
Disclaimer Above	Select a disclaimer to attach above or below emails (disclaimers created via the Mail Policies > Text Resources page or the `textconfig` command in the CLI are available in the list, see the "Text Resources" chapter.
Disclaimer Below	Select a disclaimer to attach above or below emails (disclaimers created via the Mail Policies > Text Resources page or the `textconfig` command in the CLI are available in the list, see the "Text Resources" chapter).
SMTP Authentication Profile	Specify an SMTP Authentication profile.
Certificate	Specify a certificate for TLS connections to the listener (certificates added via the Network > Certificates page or the `certconfig` command in the CLI are available in the list, see Overview of Encrypting Communication with Other MTAs).

Step 4 (Optional) Configure settings for controlling parsing in SMTP "`MAIL FROM`" and "`RCPT TO`" commands as defined in the following table.

Setting	Description
Address Parser Type	Choose how strictly the email gateway adheres to the RFC2821 standard using one of the following parser types: **Strict Mode:** • Strict mode tries to follow RFC 2821. In Strict mode, the address parser follows RFC 2821 rules with the following exceptions/enhancements: • Space is allowed after the colon, as in "MAIL FROM: <joe@example.com>". • Underscores are allowed in the domain name. • "MAIL FROM" and "RCPT TO" commands are case-insensitive. • Periods are not treated specially (for example, RFC 2821 does not allow a username of "J.D."). Some of the additional options below may be enabled which technically would violate RFC 2821. **Loose Mode:** The loose parser is basically the existing behavior from previous versions of AsyncOS. It does its best to "find" an email address and: • Ignores comments. It supports nested comments (anything found in parenthesis) and ignores them. • Does not require angle brackets around email addresses provided in "RCPT TO" and "MAIL FROM" commands. • Allows multiple nested angle brackets (it searches for the email address in the deepest nested level).
Allow 8-bit User Names	If enabled, allow 8-bit characters in the username portion of the address without escaping.
Allow 8-bit Domain Names	If enabled, allow 8-bit characters in the domain portion of the address.

Setting	Description
Allow Partial Domains	If enabled, will allow partial domains. Partial domains can be no domain at all, or a domain with no dots. The following addresses are examples of partial domains: • foo • foo@ • foo@bar This option *must* be enabled in order for the Default Domain feature to work properly. **Add Default Domain:** A default domain to use for email addresses without a fully qualified domain name. This option is disabled unless Allow Partial Domains is enabled in SMTP Address Parsing options. This affects how a listener modifies email that it relays by adding the "default sender domain" to sender and recipient addresses that do not contain fully-qualified domain names. (In other words, you can customize how a listener handles "bare" addresses). If you have a legacy system that sends email without adding (appending) your company's domain to the sender address, use this to add the default sender domain. For example, a legacy system may automatically create email that only enters the string " `joe` " as the sender of the email. Changing the default sender domain would append " `@yourdomain.com` " to " joe " to create a fully-qualified sender name of `joe@yourdomain.com` .
Source Routing	Determines behavior if source routing is detected in the "MAIL FROM" and "RCPT TO" addresses. Source routing is a special form of an email address using multiple '@' characters to specify routing (for example: @one.dom@two.dom:joe@three.dom). If set to "reject," the address will be rejected. If "strip," the source routing portion of the address will be deleted, and the message will be injected normally.
Unknown Address Literals	Determines behavior for when an address literal is received that the system cannot handle. Currently, this is everything except for IPv4. Thus, for example, for an IPv6 address literal, you can either reject it at the protocol level, or accept it and immediately hard bounce it. Recipient addresses containing literals will cause an immediate hard bounce. Sender addresses may get delivered. If the message cannot be delivered, then the hard bounce will hard bounce (double hard bounce). In the case of reject, both sender and recipient addresses will be rejected immediately at the protocol level.
Reject These Characters in User Names	Usernames that include characters (such as % or !, for example) entered here will be rejected.

Step 5 (Optional) Configure advanced settings for customizing the behavior of the listener as defined in the following table.

Setting	Description
Maximum Concurrent Connections	The maximum number of connections allowed.
TCP Listen Queue Size	The backlog of connections that AsyncOS will manage before the SMTP server accepts them.
CR and LF Handling	Choose how to handle messages that contain bare CR (carriage return) and LF (line feed) characters. • **Clean** - Allows the message but converts bare CR and LF characters to CRLF characters. Adds the **Ironport-Invalid-End-Of-Message** X-Header to all the messages with an invalid end-of-message sequence within that connection until a message with a valid end-of-message sequence is received. • **Reject** - Rejects the message. • **Allow** - Allows the message. Adds the **Ironport-Invalid-End-Of-Message** X-Header to the message with an invalid end-of-message sequence.
Add Received Header	Add a received header to all received email. A listener also modifies email that it relays by adding a Received: header on each message. If you do not want to include the Received: header, you can disable it using this option. **Note** The Received: header is not added to the message within the work queue processing. Rather, it is added when the message is enqueued for delivery Disabling the received header is a way to ensure that your network's topology is not exposed by revealing the IP addresses or hostnames of internal servers on any messages traveling outside your infrastructure. Please use caution when disabling the received header.
Use SenderBase IP Profiling	Choose whether or not to enable SenderBase IP Profiling and configure the following setting: • **SenderBase Timeout per Connection.** Define how long the appliance caches SenderBase information per SMTP connection.

Step 6 (Optional) Configure settings for controlling LDAP queries associated with this listener as defined in the following table.

Use these settings to enable LDAP queries on the listener. You must create the LDAP query first, before using this option. Each type of query has a separate subsection to configure. Click the type of query to expand the subsection.

For more information about creating LDAP queries, see LDAP Queries.

Query Type	Description
Accept Queries	For Accept queries, select the query to use from the list. You can specify whether the LDAP Accept occurs during the work queue processing or during the SMTP conversation. For LDAP Accept during the work queue processing, specify the behavior for non-matching recipients: bounce or drop. For LDAP Accept during the SMTP conversation, specify how to handle mail if the LDAP server is unreachable. You can elect to allow messages or drop the connection with a code and custom response. Finally, select whether or not to drop connections if the Directory Harvest Attack Prevention (DHAP) threshold is reached during an SMTP conversation. Performing recipient validation in the SMTP conversation can potentially reduce the latency between multiple LDAP queries. Therefore, you might notice an increased load on your directory server when you enable conversational LDAP Accept. See Overview of LDAP Queriesfor more information.
Routing Queries	For routing queries, select the query from the list. See Overview of LDAP Queries for more information.
Masquerade Queries	For masquerade queries, select a query from the list, and select which address to masquerade, such as the From or CC header addresses. See Overview of LDAP Queries for more information.
Group Queries	For group queries, select the query from the list. See Overview of LDAP Queries for more information.

Step 7 Submit and commit your changes.

What to do next

Related Topics

Partial Domains, Default Domains, and Malformed MAIL FROMs, on page 11

Partial Domains, Default Domains, and Malformed MAIL FROMs

If you enable envelope sender verification or disable allowing partial domains in SMTP Address Parsing options for a listener, the default domain settings for that listener will no longer be used.

These features are mutually exclusive.

Listening for Connection Requests by Creating a Listener Using CLI

The following table lists some of the listenerconfig subcommands used in the tasks involved in creating and editing listeners.

Table 3: Tasks for Creating Listeners

Tasks for Creating Listeners	Command(s) and Subcommands
Create a new listener	`listenerconfig -> new`
Edit global settings for listeners	`listenerconfig -> setup`
Specify a bounce profile for the listener	`bounceconfig, listenerconfig-> edit -> bounceconfig`
Associate a disclaimer with the listener	`textconfig, listenerconfig -> edit -> setup -> footer`
Configure an SMTP Authentication	`smtpauthconfig, listenerconfig -> smtpauth`
Configure SMTP address parsing	`textconfig, listenerconfig -> edit -> setup -> address`
Configure a default domain for the listener	`listenerconfig -> edit -> setup -> defaultdomain`
Add a received header to email	`listenerconfig -> edit -> setup -> received`
Change bare CR and LF characters to CRLF	`listenerconfig -> edit -> setup -> cleansmtp`
Modify the Host Access Table	`listenerconfig -> edit -> hostaccess`
Accept email for local domains or specific users (RAT) (public listeners only)	`listenerconfig -> edit -> rcptaccess`
Encrypt conversations on listeners (TLS)	`certconfig, listenerconfig -> edit`
Choose the certificate (TLS)	`listenerconfig -> edit -> certificate`

For more information about listenerconfig command, see CLI Reference Guide for AsyncOS for Cisco Secure Email Gateway.

For information about email routing and delivery configurations, see Configuring Routing and Delivery Features.

Related Topics

Advanced HAT Parameters, on page 13

Advanced HAT Parameters

The following table defines the syntax of advanced HAT parameters. Note that for the numeric values below, you can add a trailing **k** to denote kilobytes or a trailing **M** to denote megabytes. Values with no letters are considered bytes. Parameters marked with an asterisk support the variable syntax shown in the following table.

Table 4: Advanced HAT Parameter Syntax

Parameter	Syntax	Values	Example Values
Maximum messages per connection	`max_msgs_per_session`	Number	`1000`
Maximum recipients per message	`max_rcpts_per_msg`	Number	`10000` `1k`
Maximum message size	`max_message_size`	Number	`1048576` `20M`
Maximum concurrent connections allowed to this listener	`max_concurrency`	Number	`1000`
SMTP Banner Code	`smtp_banner_code`	Number	`220`
SMTP Banner Text (*)	`smtp_banner_text`	String	`Accepted`
SMTP Reject Banner Code	`smtp_banner_code`	Number	`550`
SMTP Reject Banner Text (*)	`smtp_banner_text`	String	`Rejected`
Override SMTP Banner Hostname	`use_override_hostname`	on | off | default	`default`
	`override_hostname`	String	`newhostname`
Use TLS	`tls`	on | off | required	`on`
Use anti-spam scanning	`spam_check`	on | off	`off`
Use virus scanning	`virus_check`	on | off	`off`
Maximum Recipients per Hour	`max_rcpts_per_hour`	Number	`5k`

Parameter	Syntax	Values	Example Values
Maximum Recipients per Hour Error Code	`max_rcpts_per_hour_code`	Number	`452`
Maximum Recipients per Hour Text (*)	`max_rcpts_per_hour_text`	String	`Too manyrecipients`
Use SenderBase	`use_sb`	on | off	`on`
Define IP Reputation Score	`sbrs[`*`value1`*`:`*`value2`*`]`	-10.0- 10.0	`sbrs[-10:-7.5]`
Directory Harvest Attack Prevention: Maximum Invalid Recipients Per Hour	`dhap_limit`	Number	`150`

Enterprise Gateway Configuration

In this configuration, the Enterprise Gateway configuration accepts email from the Internet and relays email to groupware servers, POP/IMAP servers, or other MTAs. At the same time, the enterprise gateway accepts SMTP messages from groupware servers and other email servers for relay to recipients on the Internet.

Figure 4: Public and Private Listeners for an Enterprise Gateway

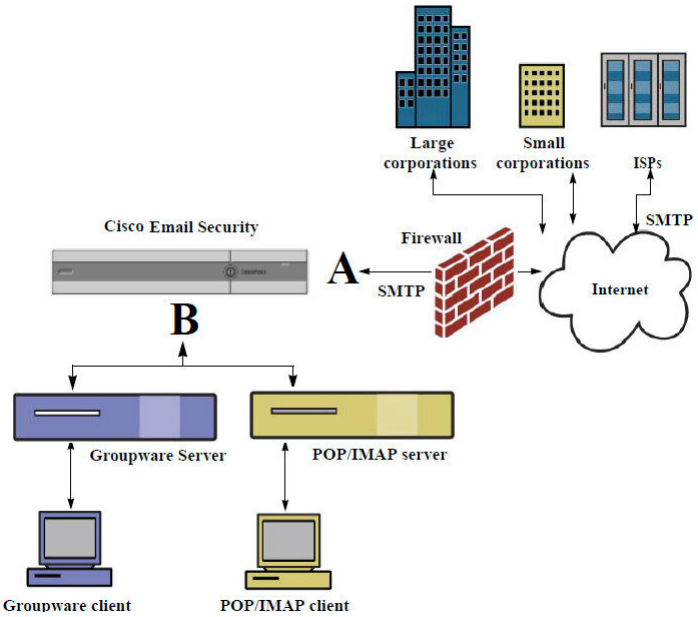

In this configuration, at least two listeners are required:

- One listener configured specifically to accept mail *from* the Internet
- One listener configured specifically to accept mail *from* your internal groupware and email servers (POP/IMAP)

By creating distinct public and private listeners for different public and private networks, you can distinguish among email for security, policy enforcement, reporting, and management. For example, email received on public listeners is scanned by your configured anti-spam engine and the anti-virus scanning engine by default, while email received on private listeners is not scanned.

Figure - Public and Private Listeners for an Enterprise Gateway shows one public listener (A) and one private listener (B) configured on the email gateway in this Enterprise Gateway configuration.