Setup and Installation

This chapter contains the following sections:

- Installation Planning, on page 1
- Physically Connecting the Email Gateway to the Network, on page 4
- Preparing for System Setup, on page 8
- Using the System Setup Wizard , on page 14
- Verifying Your Configuration and Next Steps , on page 39

Installation Planning

- Review Information That Impacts Planning Decisions , on page 1
- Plan to Place the Email Gateway at the Perimeter of Your Network, on page 1
- Register the Email Security Appliance in DNS, on page 2
- Installation Scenarios, on page 3

Review Information That Impacts Planning Decisions

- If you are configuring a virtual email gateway, please see the *Cisco Content Security Virtual Appliance Installation Guide* before continuing with this chapter.
- If you are configuring an M-Series Cisco Secure Email and Web Manager, please see Centralizing Services on a Cisco Secure Email and Web Manager (M-Series).
- We recommend reviewing Understanding the Email Pipeline before installing, as some features and functions may affect the placement of the email gateway within your infrastructure.

Plan to Place the Email Gateway at the Perimeter of Your Network

Your email gateway is designed to serve as your SMTP gateway, also known as a mail exchange (MX). For best results, some features require the email gateway to be the first machine with an IP address that is directly accessible to the Internet (that is, it is an external IP address) for sending and receiving email.

The per-recipient reputation filtering, anti-spam, anti-virus, and Virus Outbreak Filter features (see IronPort Anti-Spam Filtering, Sophos Anti-Virus Filtering, and Outbreak Filters) are designed to work with a *direct flow* of messages from the Internet and from your internal network. You can configure the email gateway for policy enforcement (Overview of Defining Which Hosts Are Allowed to Connect) for all email traffic to and from your enterprise.

Ensure that the email gateway is both accessible via the public Internet and is the "first hop" in your email infrastructure. If you allow another MTA to sit at your network's perimeter and handle all external connections, then the email gateway will not be able to determine the sender's IP address. The sender's IP address is needed to identify and distinguish senders in the Mail Flow Monitor, to query the IP Reputation Service for the sender's IP Reputation Score, and to improve the efficacy of the Anti-Spam and Outbreak Filters features.

Note If you cannot configure the email gateway as the *first* machine receiving email from the Internet, you can still exercise some of the security services available on the email gateway. For more information, see Determining Sender IP Address In Deployments with Incoming Relays.

When you use the email gateway as your SMTP gateway:

- The Mail Flow Monitor feature (see Using Email Security Monitor) offers complete visibility into all email traffic for your enterprise from both internal and external senders.
- LDAP queries (see LDAP Queries) for routing, aliasing, and masquerading can consolidate your directory infrastructure and provide for simpler updates.
- Familiar tools like alias tables (see Creating Alias Tables), domain-based routing (The Domain Map Feature), and masquerading (Configuring Masquerading) make the transition from Open-Source MTAs easier.

Register the Email Security Appliance in DNS

Malicious email senders actively search public DNS records to hunt for new victims. In order to utilize the full capabilities of Anti-Spam, Outbreak Filters, McAfee Antivirus and Sophos Anti-Virus, ensure that the email gatewayis registered in DNS.

To register the email gateway in DNS, create an A record that maps the email gateway's hostname to its IP address, and an MX record that maps your public domain to the email gateway's hostname. You must specify a priority for the MX record to advertise the email gateway as either a primary or backup MTA for your domain.

In the following example, the email gateway (ironport.example.com) is a backup MTA for the domain example.com, since its MX record has a higher priority value (20). In other words, the higher the numeric value, the lower the priority of the MTA.

```
$ host -t mx example.com

example.com mail is handled (pri=10) by mail.example.com

example.com mail is handled (pri=20) by ironport.example.com
```

By registering the email gateway in DNS, you will attract spam attacks regardless of how you set the MX record priority. However, virus attacks rarely target backup MTAs. Given this, if you want to evaluate an anti-virus engine to its fullest potential, configure the email gateway to have an MX record priority of equal or higher value than the rest of your MTAs.

Installation Scenarios

You can install your email gateway into your existing network infrastructure in several ways.

Most customers' network configurations are represented in the following scenarios. If your network configuration varies significantly and you would like assistance planning an installation, please contact Cisco Customer Support (see Cisco Customer Support).

Configuration Overview

The following figure shows the typical placement of the email gateway in an enterprise network environment:

In some scenarios, the email gateway resides inside the network "DMZ," in which case an additional firewall sits between the email gateway and the groupware server.

The following network scenarios are described:

- Behind the Firewall: two listeners configuration (*Figure - Behind the Firewall Scenario / 2 Listeners Configuration*)

Choose the configuration that best matches your infrastructure. Then proceed to the next section, Preparing for System Setup, on page 8.

Incoming

- Incoming mail is accepted for the local domains you specify.
- All other domains are rejected.
- External systems connect directly to the email gateway to transmit email for the local domains, and the email gateway relays the mail to the appropriate groupware servers (for example, Exchange™, Groupwise™, Domino™) via SMTP routes. (See Routing Email for Local Domains.)

Outgoing

- Outgoing mail sent by internal users is routed by the groupware server to the email gateway.
- The email gateway accepts outbound email based on settings in the Host Access Table for the private listener. (For more information, see Working with Listeners.)

Ethernet Interfaces

Only one of the available Ethernet interfaces on the email gateway is required in these configurations. However, you can configure two Ethernet interfaces and segregate your internal network from your external Internet network connection.

For more information about assigning multiple IP addresses to the available interfaces, see Configuring Mail Gateways for all Hosted Domains Using Virtual Gateway™ Technology and Assigning Network and IP Addresses.

Hardware Ports

The number and type of ports on your hardware appliance depend on the model:

Ports	Type	C190	C390	C690	C690F	C195	C395	C695	C695F
Management	Ethernet	0	1	1	1	0	1	1	1
Data	Ethernet	2*	5	5	3	2*	5	5	3
Console	Serial	RJ-45	RJ-45	RJ-45	RJ-45	RJ-45	RJ-45	RJ-45	RJ-45
Remote Power Management (RPC)	Ethernet	Y	Y	Y	Y	Y	Y	Y	Y

* For appliances without a dedicated management port, use the Data1 port for management purposes.

For more information about ports, see the *Hardware Installation Guide* for your appliance model.

Related Topics

- Configuring Network Interfaces, on page 19
- Accessing the Email Gateway via a Serial Connection
- Enabling Remote Power Cycling

Advanced Configurations

In addition to the configurations shown in Figure - Behind the Firewall Scenario / 2 Listeners Configuration and Figure One Listener Configuration, you can also configure:

- Multiple email gateways using the Centralized Management feature. See Centralized Management Using Clusters
- Redundancy at the network interface card level by "teaming" two of the Ethernet interfaces on email gateways using the NIC Pairing feature. See Advanced Network Configuration

Firewall Settings (NAT, Ports)

SMTP and DNS services must have access to the Internet. Other services may also require open firewall ports. For details, see Firewall Information.

Physically Connecting the Email Gateway to the Network

- Configuration Scenarios, on page 5

Configuration Scenarios

The typical configuration scenario for the email gateway is as follows:

- **Interfaces** - Only one of the three available Ethernet interfaces on the email gateway is required for most network environments. However, you can configure two Ethernet interfaces and segregate your internal network from your external Internet network connection.
- **Public Listener (incoming email)** - The public listener receives connections from many external hosts and directs messages to a limited number of internal groupware servers.

  - Accepts connections from external mail hosts based on settings in the Host Access Table (HAT). By default, the HAT is configured to ACCEPT connections from all external mail hosts.
  - Accepts incoming mail only if it is addressed for the local domains specified in the Recipient Access Table (RAT). All other domains are rejected.
  - Relays mail to the appropriate internal groupware server, as defined by SMTP Routes.

- **Private Listener (outgoing email)** - The private listener receives connections from a limited number of internal groupware servers and directs messages to many external mail hosts.

  - Internal groupware servers are configured to route outgoing mail to the Cisco C- or X-Series email gateway.
  - The email gateway accepts connections from internal groupware servers based on settings in the HAT. By default, the HAT is configured to RELAY connections from all internal mail hosts.

Related Topics

- Segregating Incoming and Outgoing Mail, on page 5

Segregating Incoming and Outgoing Mail

You can segregate incoming and outgoing email traffic over separate listeners and on separate IP addresses. You can use Internet Protocol version 4 (IPv4) and version 6 (IPv6) addresses. However, the System Setup Wizard on the email gateway supports initial configuration of the following configurations:

- 2 separate listeners on 2 logical IPv4 and 2 IPv6 addresses configured on *separate* physical interfaces

  - segregates incoming and outgoing traffic
  - you can assign an IPv4 and an IPv6 address to each listener

- 1 listener on 1 logical IPv4 address configured on one physical interface

  - combines both incoming and outgoing traffic
  - you can assign both an IPv4 and an IPv6 address to the listener

Configuration worksheets for both one and two listener configurations are included below (see Gathering the Setup Information, on page 11). Most configuration scenarios are represented by one of the following three figures.

Figure 1: Behind the Firewall Scenario / 2 Listeners Configuration
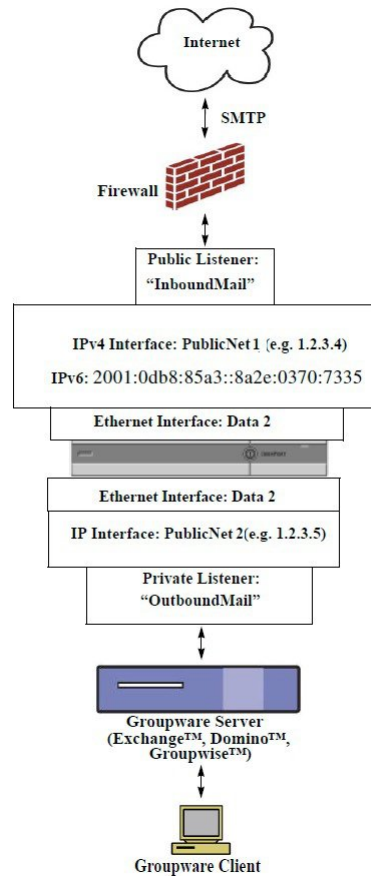
Notes:

- 2 Listeners

- 2 IPv4 addresses

- 2 IPv6 addresses

- 1 or 2 Ethernet interfaces (only 1 interface shown)
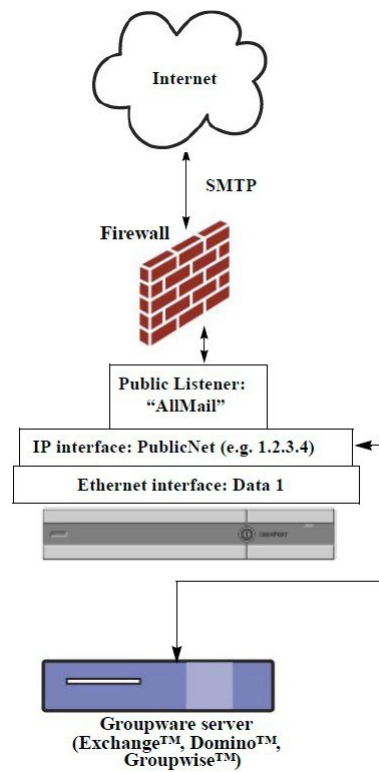
- SMTP routes configured

Inbound Listener: "InboundMail" (public)

- IPv4 address: 1.2.3.4

- IPv6 address: 2001:0db8:85a3::8a2e:0370:7334

- Listener on the Data2 interface listens on port 25

- HAT (accept ALL)

- RAT (accept mail for local domains; reject ALL)

Outbound Listener: "OutboundMail" (private)

- IP address: 1.2.3.5

- IPv6 address: 2001:0db8:85a3::8a2e:0370:7335

- Listener on the Data2 interface listens on port 25

- HAT (relay for local domains; reject ALL)

DNS can be configured to use Internet Root servers or internal DNS servers

SMTP routes direct mail to proper groupware server

Firewall ports opened for appropriate services to and from the email gateway

Figure 2: One Listener Configuration

Notes:

- 1 Listener

- 1 IP addresses

- 1 Ethernet interface

- SMTP routes configured

Inbound Listener: "InboundMail" (public)

- IP address: 1.2.3.4

- Listener on the Data2 interface listens on port 25

• HAT (accept ALL) includes entries for Groupware servers in RELAYLIST

• RAT (accept mail for local domains; reject ALL)

DNS can be configured to use Internet Root servers or internal DNS servers

SMTP routes direct mail to proper groupware server

Firewall ports opened for appropriate services to and from the email gateway.

Preparing for System Setup

• Determine Method for Connecting to the Email Gateway, on page 9
• Determining Network and IP Address Assignments, on page 10
• Gathering the Setup Information, on page 11

Procedure

	Command or Action	Purpose
Step 1	Determine how you will connect to the email gateway.	See Determine Method for Connecting to the Email Gateway, on page 9
Step 2	Determine network and IP address assignments. • If you have already cabled your email gateway to your network, ensure that the default IP address for the email gatewaydoes not conflict with other IP addresses on your network.	See Determine Method for Connecting to the Email Gateway, on page 9 and Determining Network and IP Address Assignments, on page 10
Step 3	Gather information about your system setup.	See Gathering the Setup Information, on page 11.
Step 4	Review the latest product release notes for your email gateway.	Release notes are available from the link in Documentation.
Step 5	Unpack the email gateway, physically install it in a rack, and turn it on.	See Quickstart Guide for your email gateway. This guide is available from the link in Documentation.
Step 6	If you will run the setup wizard using the command line interface (CLI), access the CLI.	See Running the Command Line Interface (CLI) System Setup Wizard, on page 26)
Step 7	If you will run the setup wizard using the web interface:	**a.** (Virtual appliances ONLY) Access the command-line interface and enable HTTP and/or HTTPS using the `interfaceconfig` command. **b.** Launch a web browser and enter the IP address of the email gateway.
Step 8	If you are setting up a virtual email gateway, load your virtual email gateway license.	Use the `loadlicense` command. For more information, see the *Cisco Content Security*

	Command or Action	Purpose
		Virtual Appliance Installation Guide available from the link in Documentation.
Step 9	Configure basic settings for your system.	See Using the System Setup Wizard , on page 14

Determine Method for Connecting to the Email Gateway

To successfully set up the email gateway in your environment, you must gather important network information from your network administrator about how you would like to connect the email gateway to your network.

Related Topics

- Connecting to the Email Gateway, on page 9

Connecting to the Email Gateway

During the initial setup, you can connect to the email gateway in one of two ways:

Table 1: Options for Connecting to the Email Gateway

Ethernet	An Ethernet connection between a PC and the network and between the network and the Management port. The IPv4 address that has been assigned to the Management port by the factory is 192.168.42.42 . This is the easiest way to connect if it works with your network configuration.
Serial	A serial communications connection between the PC and the Serial Console port. If you cannot use the Ethernet method, a straight serial-to- serial connection between the computer and the email gateway will work until alternate network settings can be applied to the Management port. For pinout information, see Accessing the Email Gateway via a Serial Connection. The communications settings for the serial port are: **Bits per second:** 9600 **Data bits:** 8 **Parity:** None **Stop bits:** 1 **Flow control:** Hardware

Note Keep in mind that the initial connection method is not final. This process applies only for the initial configuration. You can change network settings at a later time to allow different connection methods. (See FTP, SSH, and SCP Access for more information.) You can also create multiple user accounts with differing administrative privileges to access the email gateway. (For more information, see Adding Users.)

Determining Network and IP Address Assignments

You can use both IPv4 and IPv6 addresses.

Default IP Addresses for Management and Data Ports

The IP address that is pre-configured on the Management port (the Data 1 port on C170 and C190 email gateways) is 192.168.42.42 .

Choosing Network Connections to Receive and Deliver Email

Most users take advantage of the two Data Ethernet ports on the email gateway by connecting to two networks from the email gateway:

- The private network accepts and delivers messages to your internal systems.
- The public network accepts and delivers messages to the Internet.

Other users may want to use only one Data port serving both functions. Although the Management Ethernet port can support any function, it is preconfigured for access to the graphical user interface and the command line interface.

Binding Logical IP Addresses to Physical Ethernet Ports

You can segregate incoming and outgoing email traffic over separate listeners and on separate IP addresses. You can use Internet Protocol version 4 (IPv4) and version 6 (IPv6) addresses. However, the System Setup Wizard on the email gateway supports initial configuration of the following configurations:

- 2 separate listeners on 2 logical IPv4 and 2 IPv6 addresses configured on *separate* physical interfaces

  - segregates incoming and outgoing traffic
  - you can assign an IPv4 and an IPv6 address to each listener

- 1 listener on 1 logical IPv4 address configured on one physical interface

  - combines both incoming and outgoing traffic
  - you can assign both an IPv4 and an IPv6 address to the listener

The email gateway can support both IPv4 and IPv6 addresses on single listener. The listener will accept mail on both the addresses. All settings on a listener apply to both IPv4 and IPv6 addresses.

Choosing Network Settings for Your Connections

You will need the following network information about each Ethernet port that you choose to use:

- IP address (IPv4 or IPv6 or both)
- Netmask for IPv4 address in CIDR format
- Prefix for IPv6 address in CIDR format

In addition, you will need the following information about your overall network:

- IP address of the default router (gateway) on your network

- IP address and hostname of your DNS servers (not required if you want to use Internet root servers)
- Hostname or IP address of your NTP servers (not required if you want to use Cisco's time servers)

See Assigning Network and IP Addresses for more information.

Note If you are running a firewall on your network between the Internet and the email gateway, it may be necessary to open specific ports for the email gateway to work properly. See Firewall Information for more information.

Gathering the Setup Information

Now that you understand the requirements and strategies when making the necessary selections in the System Setup Wizard, use the following tables to gather information about your system setup while reading this section.

See Assigning Network and IP Addresses for more detailed information on network and IP addresses. See Centralizing Services on a Cisco Secure Email and Web Manager (M-Series) if you are configuring a Cisco Secure Email and Web Manager.

Table 2: System Setup Worksheet: 2 Listeners for Segregating Email Traffic

System Settings		
Default System Hostname:		
Email System Alerts To:		
Deliver Scheduled Reports To:		
Time Zone Information:		
NTP Server:		
Admin Passphrase:		
SenderBase Network Participation:	Enable / Disable	
AutoSupport:	Enable / Disable	
Network Integration		
Gateway:		
DNS (Internet or Specify Own):		
Interfaces		
Data 1 Port		
IPv4 Address / Netmask:		
IPv6 Address / Prefix:		

System Settings		
Fully Qualified Hostname:		
Accept Incoming Mail:	**Domain**	**Destination**
Relay Outgoing Mail:	**System**	
Data 2 Port		
IPv4 Address / Netmask:		
IPv6 Address / Prefix:		
Fully Qualified Hostname:		
Accept Incoming Mail:	**Domain**	**Destination**
Relay Outgoing Mail:	**System**	
Management Port		
IP Address:		
Network Mask:		
IPv6 Address:		
Prefix:		
Fully Qualified Hostname:		
Accept Incoming Mail:	**Domain**	**Destination**
Relay Outgoing Mail:	**System**	
Message Security		
IP Reputation Filtering:	Enable / Disable	
Anti-Spam Scanning Engine	None / IronPort	
McAfee Anti-Virus Scanning Engine	Enable / Disable	
Sophos Anti-Virus Scanning Engine	Enable / Disable	
Outbreak Filters	Enable / Disable	

Table 3: System Setup Worksheet: 1 Listener for All Email Traffic

System Settings		
Default System Hostname:		

System Settings		
Email System Alerts To:		
Deliver Scheduled Reports To:		
Time Zone:		
NTP Server:		
Admin Passphrase:		
SenderBase Network Participation:	Enable / Disable	
AutoSupport:	Enable / Disable	
Network Integration		
Gateway:		
DNS (Internet or Specify Own):		
Interfaces		
Data2 Port		
IPv4 Address / Netmask:		
IPv6 Address / Prefix:		
Fully Qualified Hostname:		
Accept Incoming Mail:	**Domain**	**Destination**
Relay Outgoing Mail:	**System**	
Data1 Port		
IPv4 Address / Netmask:		
IPv6 Address / Prefix:		
Fully Qualified Hostname:		
Message Security		
IP Reputation Filtering:	Enable / Disable	
Anti-Spam Scanning Engine	None / IronPort	
McAfee Anti-Virus Scanning Engine	Enable / Disable	
Sophos Anti-Virus Scanning Engine	Enable / Disable	

System Settings		
Outbreak Filters	Enable / Disable	

Using the System Setup Wizard

You must use the System Setup Wizard for the initial setup in order to ensure a complete configuration. Later, you can configure custom options not available in the System Setup Wizard.

You can run the System Setup Wizard using a browser or the command line interface (CLI). For more information, see Accessing the Web-Based Graphical User Interface (GUI), on page 14 or Running the Command Line Interface (CLI) System Setup Wizard, on page 26

Before you begin, complete the prerequisites at Preparing for System Setup, on page 8.

⚠
Caution   If you are setting up a virtual email gateway, you will have to use the loadlicense command to load your virtual email gateway license before running the System Setup Wizard. See the *Cisco Content Security Virtual Appliance Installation Guide* for more information.

⚠
Caution   The System Setup Wizard will completely reconfigure your system. You should only use the System Setup Wizard the very first time you install the email gateway, or if you want to completely overwrite your existing configuration.

⚠
Caution   The email gateway ships with a default IP address of 192.168.42.42 on the Management port of all hardware except C170 and C190 email gateways, which use the Data 1 port instead. Before connecting the email gateway to your network, ensure that no other device's IP address conflicts with this factory default setting. If you are configuring a Cisco Secure Email and Web Manager, see Centralizing Services on a Cisco Secure Email and Web Manager (M-Series).

If you are connecting multiple factory-configured content security appliances to your network, add them one at a time, reconfiguring each appliance's default IP address as you go.

Accessing the Web-Based Graphical User Interface (GUI)

The email gateway has a standard web-based graphical user interface, a new web-based interface for managing the Email Security Monitor feature (Monitoring, Tracking, and Quarantine), and a command-line interface.

To access the web-based Graphical User Interface (GUI), open your web browser and point it to 192.168.42.42.

[New Web Interface Only] You can access the new web interface in any one of the following ways:

✎

Note The new web interface of your email gateway uses AsyncOS API HTTP/HTTPS ports (6080/6443) and trailblazer HTTPS port (4431). You can use the `trailblazerconfig` command in the CLI to configure the trailblazer HTTPS ports. Make sure that the trailblazer HTTPS port is opened on the firewall.

- When `trailblazerconfig` CLI command is enabled, use the following URL - `https://example.com:<trailblazer-https-port>/ng-login`

  where `example.com` is the email gateway host name and `<trailblazer-https-port>` is the trailblazer HTTPS port configured on the email gateway.

  For more information on the trailblazerconfig CLI command, see the Cisco Secure Email Command Reference Guide.

- Log in to the legacy web interface and click **Secure Email Gateway is getting a new look. Try it!!** link to access the new web interface.

Important Notes

- Make sure that AsyncOS API is enabled on the email gateway.

- Make sure that AsyncOS HTTPS API port is not enabled on multiple interfaces.

- You must login to the legacy web interface of the email gateway.

- If `trailblazerconfig` is enabled, the configured HTTPS port must be opened on the firewall. The default HTTPS port is 4431.

  Also ensure that your DNS server can resolve the hostname that you specified for accessing the email gateway.

Related Topics

Factory Default Username and Passphrase

If you install a new virtual or hardware email gateway, you must change the default passphrase to get complete access to set up the email gateway. When you log in to the email gateway for the first time, the web interface prompts you to change the default passphrase, and the CLI limits the access to the following commands till you change the default passphrase.

- Commit

- Interfaceconfig

- passphrase

- Loadconfig

- Systemsetup

• loadlicense (for virtual email gateways)

• feature key

• Ping

• Telnet

• netstat

• Username: **admin**
• Passphrase: **ironport**

For Example:

```
login: admin
passphrase: ironport
```

Note If your session times out, you will be asked to re-enter your username and passphrase. If your session times out while you are running the System Setup Wizard, you will have to start over again.

Accessing the Legacy Web Interface

To access the legacy web interface from the new web interface, click on the gear icon ⚙ as shown in the following figure:

Figure 3: Accessing the Legacy Web Interface from the

| Monitoring | Tracking | Quarantine | admin 👤 ? ⚙ |

The legacy web interface opens in a new browser window. You must log in again to access it.

If you want to log out of the appliance completely, you need to log out of both the new and legacy web interfaces of your email gateway.

Accessing New Web Interface in Dusk Mode

Dusk Mode is a reversed color scheme that utilizes light-colored typography, UI elements, and iconography on dark backgrounds.

You can now access the new web interface of your email gateway using the dusk mode.

To switch to the dusk mode, click on the user icon on the top-right corner of the new web interface and select **Dusk Theme**.

Defining Basic Configuration Using the Web-Based System Setup Wizard

Procedure

Step 1 Launch the System Setup Wizard

- Log in to the graphical user interface as described in Accessing the Web-Based Graphical User Interface (GUI), on page 14.

- On brand new (not upgraded from previous releases of AsyncOS) systems, your browser will automatically be redirected to the System Setup Wizard.

- Otherwise, on the System Administration tab, click System Setup Wizard in the list of links on the left.

Step 2   Start. See Step 1: Start, on page 17.

- Read and accept the license agreement

Step 3   System. See Step 2: System, on page 18.

- Setting the hostname of the email gateway
- Configuring alert settings, report delivery settings, and AutoSupport
- Setting the system time settings, and NTP server
- Resetting the admin passphrase
- Enabling Service Logs

Step 4   Network. See Step 3: Network, on page 19.

- Defining the default router and DNS settings
- Enabling and configuring network interfaces, including:Configuring incoming mail (inbound listener)Defining SMTP routes (optional) Configuring outgoing mail (outbound listener) and defining systems allowed to relay mail through the email gateway (optional)

Step 5   Security. Se Step 4: Security, on page 23.

- Enabling IP Reputation Filtering
- Enabling the Anti-Spam service
- Enabling the Spam Quarantine
- Enabling the Anti-Virus service
- Enabling Advanced Malware Protection (file reputation and analysis services.)
- Enabling the Outbreak Filters service

Step 6   Review. See Step 5: Review, on page 24.

- Reviewing your setup and installing the configuration
- At the end of the process, you are prompted to

Step 7   Commit the changes you have made.

Your changes will not take effect until they have been committed.

Step 1: Start

Begin by reading the license agreement. Once you have read and agreed to the license agreement, check the box indicating that you agree and then click **Begin Setup** to proceed.

You can also view the text of the agreement here: https://support.ironport.com/license/eula.html

Step 2: System

Setting the Hostname

Define the fully-qualified hostname for the email gateway. This name should be assigned by your network administrator.

Configuring System Alerts

Cisco AsyncOS sends alert messages via email if there is a system error that requires the user's intervention. Enter the email address (or addresses) to which to send those alerts.

You must add at least one email address that receives System Alerts. Enter a single email address, or separate multiple addresses with commas. The email recipients initially receive all types of alerts at all levels, except for Directory Harvest Attack Prevention alerts. You can add more granularity to the alert configuration later. For more information, see Alerts.

Configuring Report Delivery

Enter the address to which to send the default scheduled reports. If you leave this value blank, the scheduled reports are still run. They will be archived on the email gateway rather than delivered.

Setting the Time

Set the time zone on the email gateway so that timestamps in message headers and log files are correct. Use the drop-down menus to locate your time zone or to define the time zone via GMT offset (see Selecting a GMT Offset for more information).

You can set the system clock time manually later, or you can use the Network Time Protocol (NTP) to synchronize time with other servers on your network or the Internet. By default, one entry to the Cisco Systems time servers (`time.ironport.com`) to synchronize the time on your email gateway is already configured.

Setting the Passphrase

Set the passphrase for the admin account. This is a required step. When changing the passphrase for the Cisco AsyncOS admin account, the new passphrase must be six characters or longer. Be sure to keep the passphrase in a secure location.

In addition to creating a login passphrase manually, you can also create a system-generated passphrase to log in to your email gateway.

Enabling Service Logs

The 'Service Logs' are sent to the Cisco Talos Cloud service to improve Phishing detection.

If you enable Service Logs, the Cisco Email Security gateway collects limited personal data from customer emails and offers extensive useful threat detection capabilities that can be coupled with dedicated analysis systems to collect, trend, and correlate observed threat activity. Cisco uses the personal data to improve your

email gateway capabilities to analyze the threat landscape, provide threat classification solutions on malicious emails, and to protect your email gateway from new threats such as spam, virus, and directory harvest attacks.

See Improving Phishing Detection Efficacy using Service Logs for more information.

Enabling AutoSupport

The AutoSupport feature (enabled by default) keeps the Cisco Customer Support team aware of issues with your email gateway so that we can provide better support to you. (For more information, see AutoSupport.)

Click **Next** to continue.

Step 3: Network

In Step 3, you define the default router (gateway) and configure the DNS settings, and then set up the email gateway to receive and or relay email by configuring the Data 1, Data 2, and Management interfaces.

Configuring DNS and Default Gateway

Type the IP address of the default router (gateway) on your network. You can use an IPv4 address, an IPv6 address, or both.

Next, configure the DNS (Domain Name Service) settings. Cisco AsyncOS contains a high-performance internal DNS resolver/cache that can query the Internet's root servers directly, or the system can use DNS servers you specify. If you choose to use your own servers, you will need to supply the IP address and hostname of each DNS server. You can enter up to four DNS servers via the System Setup Wizard. Please note that DNS servers you enter will have an initial priority of 0. For more information, see Configuring Domain Name System (DNS) Settings.

Note   The email gateway requires access to a working DNS server in order to perform DNS lookups for incoming connections. If you cannot specify a working DNS server that is reachable by the email gateway while you are setting up the email gateway, a workaround is to either select "Use Internet Root DNS Servers" or to specify, temporarily, the IP address of the Management interface so that you can complete the System Setup Wizard.

Configuring Network Interfaces

Your email gateway has network interfaces that are associated with the physical Ethernet ports on the machine.

To use an interface, mark the "Enable" checkbox and then specify an IP address, network mask, and fully qualified hostname. The IP address you enter should be the address intended for your inbound mail as reflected in your DNS records. Typically this address would have an MX record associated with it in DNS. You can use an IPv4 address, an IPv6 address, or both. If you use both, the interface will accept both types of connections.

Each interface can be configured to accept mail (incoming), relay email (outgoing), or appliance management. During setup, you are limited to one of each. On most email gateways, you would typically use one interface

for incoming, one for outgoing, and one for appliance management. On the C170 and C190 email gateways, you would typically use one interface for both incoming and outgoing mail, and the other interface for management.

You must configure one interface to receive email.

Assign and configure a logical IP address to one of the physical Ethernet interfaces on the email gateway. If you decide to use both the Data 1 Ethernet port and the Data 2 Ethernet port, you need this information for both connections.

For C390, and C690 appliances: Cisco recommends using one of the physical Ethernet ports to connect directly to the Internet for the purposes of receiving inbound email through public listeners, and using another physical Ethernet port to connect directly to your internal network for the purposes of relaying outbound email through private listeners.

For C190 appliances: Typically, the System Setup Wizard will configure only one physical Ethernet port with one listener for both receiving inbound email and relaying outbound email.

See Binding Logical IP Addresses to Physical Ethernet Ports, on page 10.

The following information is required:

- The **IP address** assigned by your network administrator. This can be an IPv4 address, an IPv6 address, or both.
- For IPv4 addresses: the **netmask** of the interface. AsyncOS only accepts a netmask in CIDR format. For example, `/24` for the `255.255.255.0` subnet.

  For IPv6 addresses: the **prefix** in CIDR format. For example `/64` for a 64-bit prefix.

- (optional) A fully-qualified hostname for the IP address.

Note   IP addresses within the same subnet cannot be configured on separate physical Ethernet interfaces. See Assigning Network and IP Addresses for more detailed information on Network and IP Address configuration.

Accepting Mail

When configuring your interfaces to accept mail, you define:

- the domain for which to accept mail
- destination (SMTP Route) for each domain, this is optional

Mark the checkbox for Accept Incoming Mail to configure the interface to accept mail. Enter the name of the domain for which to accept mail.

Enter the Destination. This is the SMTP Route or name of the machine(s) where you would like to route email for the domains specified.

This is the first SMTP Routes entry. The SMTP Routes table allows you to redirect all email for each domain (also known as a Recipient Access Table (RAT) entry) you enter to a specific mail exchange (MX) host. In typical installations, the SMTP Routes table defines the specific groupware (for example, Microsoft Exchange) server or the "next hop" in the email delivery for your infrastructure.

For example, you can define a route that specifies that mail accepted for the domain `example.com` and all of its subdomains `.example.com` is routed the to the groupware server `exchange.example.com`.

You can enter multiple domains and destinations. Click **Add Row** to add another domain. Click the trash can icon to remove a row.

✎

Note Configuring SMTP Routes in this step is optional. If no SMTP routes are defined, the system will use DNS to lookup and determine the delivery host for the incoming mail received by the listener. (See Routing Email for Local Domains.)

You must add at least one domain to the Recipient Access Table. Enter a domain — `example.com`, for example. To ensure that mail destined for any subdomain of `example.net` will match in the Recipient Access Table, enter `.example.net` as well as the domain name. For more information, see Defining Recipient Addresses.

Relaying Mail (Optional)

When configuring your interfaces to relay mail, you define the systems allowed to relay email through the email gateway.

These are entries in the RELAYLIST of the Host Access Table for a listener. See Sender Group Syntax for more information.

Mark the check box for Relay Outgoing Mail to configure the interface to relay mail. Enter the hosts that may relay mail through the email gateway.

When you configure an interface to relay outbound mail, the System Setup Wizard turns on SSH for the interface as long as no public listeners are configured to use the interface.

In the following example, two interfaces with IPv4 addresses are created:

- 192.168.42.42 remains configured on the Management interface.
- 192.168.1.1 is enabled on the Data 1 Ethernet interface. It is configured to accept mail for domains ending in .example.com and an SMTP route is defined for exchange.example.com.
- 192.168.2.1 is enabled on the Data 2 Ethernet interface. It is configured to relay mail from exchange.example.com.

C390, and C690 Installations

Figure 4: Network Interfaces: 2 Interfaces in Addition to Management (Segregated Traffic)

C170 and C190 Installations

For C170 and C190 appliances, the Data 2 interface is typically configured for both incoming and outgoing mail while the Data 1 interface is used for appliance management.

When configuring a single IP address for all email traffic (nonsegregated traffic), step 3 of the System Setup Wizard will look like this:

Figure 5: Network Interfaces: 1 IP Address for Incoming and Outgoing (Nonsegregated) Traffic

Click **Next** to continue.

Step 4: Security

In step 4, you configure anti-spam and anti-virus settings. The anti-spam options include IP Reputation Filtering and selecting an anti-spam scanning engine. For anti-virus, you can enable Outbreak Filters and Sophos or McAfee anti-virus scanning.

- Enabling IP Reputation Filtering, on page 23
- Enabling Anti-Spam Scanning, on page 23
- Enabling Anti-Virus Scanning, on page 23
- Enabling Advanced Malware Protection (File Reputation and Analysis Services) , on page 23
- Enabling Outbreak Filters, on page 24

Enabling IP Reputation Filtering

The IP Reputation Service can be used as a stand-alone anti-spam solution, but it is primarily designed to improve the effectiveness of a content-based anti-spam system such as Anti-Spam.

The IP Reputation Service provides an accurate, flexible way for users to reject or throttle suspected spam based on the connecting IP address of the remote host. The IP Reputation Service returns a score based on the probability that a message from a given source is spam. The IP Reputation Service is unique in that it provides a global view of email message volume and organizes the data in a way that makes it easy to identify and group related sources of email. Cisco strongly suggests that you enable IP Reputation Filtering.

Once enabled, IP Reputation Filtering is applied on the incoming (accepting) listener.

Enabling Anti-Spam Scanning

Your email gateway may ship with a 30-day evaluation key for Anti-Spam software. During this portion of the System Setup Wizard, you can choose to enable Anti-Spam globally on the email gateway. You can also elect to not enable the service.

If you choose to enable the anti-spam service, you can configure AsyncOS to send spam and suspected spam messages to the local Spam Quarantine. The Spam Quarantine serves as the end-user quarantine for the email gateway. Only administrators can access the quarantine until end-user access is configured.

See Managing Spam and Graymail for all of the Anti-Spam configuration options available on the email gateway. See Policy, Virus, and Outbreak Quarantines.

Enabling Anti-Virus Scanning

Your email gateway may ship with a 30-day evaluation key for the Sophos Anti-Virus or McAfee Anti-Virus scanning engines. During this portion of the System Setup Wizard, you can choose to enable an anti-virus scanning engine globally on the email gateway.

If you choose to enable an anti-virus scanning engine, it is enabled for both the default incoming and default outgoing mail policies. The email gateway scans mail for viruses, but it does not repair infected attachments. The email gateway drops infected messages.

See Anti-Virus for all of the anti-virus configuration options available on the email gateway.

Enabling Advanced Malware Protection (File Reputation and Analysis Services)

Advanced Malware Protection obtains reputation information about attached files from a cloud-based service.

For more information, see File Reputation Filtering and File Analysis

Enabling Outbreak Filters

Your email gateway may ship with a 30-day evaluation key for Outbreak Filters. Outbreak Filters provide a "first line of defense" against new virus outbreaks by quarantining suspicious messages until traditional anti-virus security services can be updated with a new virus signature file.

See Outbreak Filters for more information.

Click **Next** to continue.

Step 5: Review

A summary of the configuration information is displayed. You can edit the System Settings, Network Integration, and Message Security information by clicking the **Previous** button or by clicking the corresponding **Edit** link in the upper-right of each section. When you return to a step to make a change, you must proceed through the remaining steps until you reach this review page again. All settings you previously entered will be remembered.

Once you are satisfied with the information displayed click **Install This Configuration**.

A confirmation dialog is displayed. Click **Install** to install the new configuration.

Your email gateway is now ready to send email.

Note   Clicking **Install** will cause the connection to the current URL (http://192.168.42.42) to be lost if you changed the IP address of the interface you used to connect to the email gateway from the default. However, your browser will be redirected to the new IP address.

Once System Setup is complete, several alert messages are sent. See Immediate Alerts, on page 39 for more information.

Setting up the Connection to Active Directory

If the System Setup Wizard properly installs the configuration on the email gateway, the Active Directory Wizard appears. If you are running an Active Directory server on your network, use the Active Directory Wizard to configure an LDAP server profile for the Active Directory server and assign a listener for recipient validation. If you are not using Active Directory or want to configure it later, click Skip this Step. You can run the Active Directory Wizard on the **System Administration > Active Directory Wizard** page. You can also configure Active Directory and other LDAP profiles on the **System Administration > LDAP** page.

The Active Directory Wizard retrieves the system information needed to create an LDAP server profile, such as the authentication method, the port, the base DN, and whether SSL is supported. The Active Directory Wizard also creates LDAP accept and group queries for the LDAP server profile.

After the Active Directory Wizard creates the LDAP server profile, use the **System Administration > LDAP** page to view the new profile and make additional changes . It is recommended that you avoid changing LDAP settings on Cisco Secure Email Cloud Gateway.

Procedure

Step 1   On the Active Directory Wizard page, click **Run Active Directory Wizard**.

Step 2   Enter the host name for the Active Directory server.

Step 3   Enter a username and passphrase for the authentication request.

Step 4   Click **Next** to continue.

              The Active Directory Wizard tests the connection to the Active Directory server. If successful, the Test Directory Settings page is displayed.

Step 5   Test the directory settings by entering an email address that you know exists in the Active Directory and clicking **Test**. The results appear in the connection status field.

Step 6   Click **Done**.

Proceeding to the Next Steps

After you successfully configure your email gateway to work with your Active Directory Wizard, or skip the process, the System Setup Next Steps page appears.

Click the links on the System Setup Next Steps page to proceed with the configuration of your email gateway.

Accessing the Command Line Interface (CLI)

Access to the CLI varies depending on the management connection method you chose in Connecting to the Email Gateway, on page 9. The factory default username and passphrase are listed next. Initially, only the admin user account has access to the CLI. You can add other users with differing levels of permission after you have accessed the command line interface for the first time via the admin account. (For information about adding users, see Adding Users.) The System Setup Wizard asks you to change the passphrase for the admin account. The passphrase for the admin account can also be reset directly at any time using the `passphrase` command.

To connect via Ethernet: Start an SSH session with the factory default IP address 192.168.42.42. SSH is configured to use port 22. Enter the username and passphrase below.

To connect via a Serial connection: Start a terminal session with the communication port on your personal computer that the serial cable is connected to. Use the settings for serial port outlined in Connecting to the Email Gateway, on page 9. Enter the username and passphrase below.

Log in to the email gateway by entering the username and passphrase.

Related Topics

- Factory Default Username and Passphrase, on page 15

Factory Default Username and Passphrase

If you install a new virtual or hardware email gateway, you must change the default passphrase to get complete access to set up the email gateway. When you log in to the email gateway for the first time, the web interface prompts you to change the default passphrase, and the CLI limits the access to the following commands till you change the default passphrase.

- Commit

- Interfaceconfig

- passphrase

- Loadconfig

- Systemsetup

- loadlicense (for virtual email gateways)

- feature key

- Ping

- Telnet

- netstat

- Username: **admin**
- Passphrase: **ironport**

For Example:

```
login: admin
passphrase: ironport
```

Note If your session times out, you will be asked to re-enter your username and passphrase. If your session times out while you are running the System Setup Wizard, you will have to start over again.

Running the Command Line Interface (CLI) System Setup Wizard

The CLI version of the System Setup Wizard basically mirrors the steps in the GUI version, with a few minor exceptions:

- The CLI version includes prompts to enable the web interface.
- The CLI version allows you to edit the default Mail Flow Policy for each listener you create.
- The CLI version contains prompts for configuring the global Anti-Virus and Outbreak Filters security settings.
- The CLI version does not prompt you to create an LDAP profile after the system setup is complete. Use the ldapconfig command to create an LDAP profile.

To run the System Setup Wizard, type systemsetup at the command prompt.

```
IronPort> systemsetup
```

The System Setup Wizard warns you that you will reconfigure your system. If this is the very first time you are installing the email gateway, or if you want to completely overwrite your existing configuration, answer "Yes" to this question.

```
WARNING: The system setup wizard will completely delete any existing

'listeners' and all associated settings including the 'Host Access Table' -
mail operations may be interrupted.
```

```
Are you sure you wish to continue? [Y]> Y
```

✎

Note   The remainder of the system setup steps are described below. Examples of the CLI System Setup Wizard dialogue will only be included for sections that deviate from the GUI System Setup Wizard described above in Defining Basic Configuration Using the Web-Based System Setup Wizard , on page 16.

Related Topics

Change the Admin Passphrase

First, you change the passphrase for the AsyncOS admin account. You must enter the old passphrase to continue. The new passphrase must be six characters or longer. Be sure to keep the passphrase in a secure location. Changes made to the passphrase are effective once the system setup process is finished.

In addition to creating a login passphrase manually, you can also create a system-generated passphrase to log in to your email gateway.

Accept the License Agreement

Read and accept the software license agreement that is displayed.

Set the Hostname

Next, you define the fully-qualified hostname for the email gateway. This name should be assigned by your network administrator.

Assign and Configure Logical IP Interface(s)

The next step assigns and configures a logical IP interface on the physical Ethernet interface named Management (on C390, and C690 appliances) or Data 1 (on C190 appliances), and then prompts you to configure a logical IP interface on any other physical Ethernet interfaces available on the email gateway.

Each Ethernet interface can have multiple IP interfaces assigned to it. An IP interface is a logical construct that associates an IP address and hostname with a physical Ethernet interface. If you decided to use both the Data 1 and Data 2 Ethernet ports, you need the IP addresses and hostnames for both connections.

For C390, and C690 appliances: Cisco recommends using one of the physical Ethernet ports to connect directly to the Internet for the purposes of receiving inbound email through public listeners, and using another physical Ethernet port to connect directly to your internal network for the purposes of relaying outbound email through private listeners.

For C190 appliances: By default, the `systemsetup` command will configure only one physical Ethernet port with one listener for receiving inbound email and relaying outbound email.

> **Note** When you configure an interface to relay outbound mail, the system turns on SSH for the interface as long as no public listeners are configured to use the interface.

The following information is required:

- A **name** (nickname) created by you to refer to the IP interface later. For example, if you are using one Ethernet port for your private network and the other for the public network, you may want to name them PrivateNet and PublicNet, respectively.

  > **Note** The names you define for interfaces are case-sensitive. AsyncOS will not allow you to create two identical interface names. For example, the names `Privatenet` and `PrivateNet` are considered as two *different* (unique) names.

- The IP ad**dress** assigned by your network administrator. This is can be an IPv4 or IPv6 address, You can assign both types of IP addresses to a single IP interface.
- The **netmask** of the interface. The netmask must be in CIDR format. For example, use `/24` for the `255.255.255.0` subnet.

  > **Note** IP addresses within the same subnet cannot be configured on separate physical Ethernet interfaces. See Assigning Network and IP Addresses for more detailed information on Network and IP Address configuration.
  >
  > For C190 appliances, the Data 2 interface is configured first.

Specify the Default Gateway

In the next portion of the `systemsetup` command, you type the IP address of the default router (gateway) on your network.

Enable the Web Interface

In the next portion of the `systemsetup` command, you enable the web interface for the email gateway (for the Management Ethernet interface). You can also choose to run the web interface over secure HTTP (https). If you choose to use HTTPS, the system will use a demonstration certificate until you upload your own certificate.

Configure the DNS Settings

Next, you configure the DNS (Domain Name Service) settings. Cisco AsyncOS contains a high-performance internal DNS resolver/cache that can query the Internet's root servers directly, or the system can use your own DNS servers. If you choose to use your own servers, you will need to supply the IP address and hostname of each DNS server. You can enter as many DNS servers as you need (each server will have a priority of 0.). By default, `systemsetup` prompts you to enter the addresses for your own DNS servers.

Create a Listener

A "listener" manages inbound email processing services that will be configured on a particular IP interface. Listeners only apply to email entering the email gateway — either from your internal systems or from the Internet. Cisco AsyncOS uses listeners to specify criteria that messages must meet in order to be accepted and relayed to recipient hosts. You can think of a listener as an email listener (or even a "SMTP daemon") running for IP addresses you specified above.

For C390, and C690 appliances: By default, the `systemsetup` command configures two listeners — one public and one private. (For more information on the types of listeners available, see Configuring the Gateway to Receive Email.)

For C190 appliances: By default, the `systemsetup` command configures one public listener for both receiving mail from the Internet and for relaying email from your internal network. See Listener Example for C190 Appliances , on page 34.

When you define a listener, you specify the following attributes:

- A **name** (nickname) created by you to refer to the listener later. For example, the listener that accepts email from your internal systems to be delivered to the Internet may be called OutboundMail.
- One of the IP interfaces (that you created earlier in the `systemsetup` command) on which to receive email.
- The name of the machine(s) to which you want to route email (public listeners only). (This is the first `smtproutes` entry. See Routing Email for Local Domains.)
- Whether or not to enable filtering based on IP Reputation Scores for public listeners. If enabled, you are also prompted to select between Conservative, Moderate, or Aggressive settings.
- Rate-limiting per host: the maximum number of recipients per hour you are willing to receive from a remote host (public listeners only).
- The recipient domains or specific addresses you want to accept email for (public listeners) or the systems allowed to relay email through the email gateway (private listeners). (These are the first Recipient Access Table and Host Access Table entries for a listener. See Sender Group Syntax and Adding Domains and Users For Which to Accept Messages for more information.)

Related Topics

Public Listener

> ✎
>
> **Note** The following examples of creating a public and private listener apply to C390, and C690 appliances only. For C190 appliances, skip to the next section, Listener Example for C190 Appliances , on page 34.

In this example portion of the `systemsetup` command, a public listener named InboundMail is configured to run on the PublicNet IP interface. Then, it is configured to accept all email for the domain `example.com` . An initial SMTP route to the mail exchange `exchange.example.com` is configured. Rate limiting is enabled, and the maximum value of 4500 recipients per hour from a single host is specified for the public listener.

> ✎
>
> **Note** The value you enter for maximum recipients per hour you are willing to receive from a remote host is a completely arbitrary value, one that is usually relative to the size of the enterprise for which you are administering email. For example, a sender who sends 200 messages in an hour might be considered a "spammer" (sender of unsolicited bulk email), but if you are configuring the email gateway to handle all email for a 10,000 person company, 200 messages per hour from a remote host may be a reasonable value. Conversely, in a 50-person company, someone sending 200 messages in an hour to you may be an obvious spammer. You must choose an appropriate value when you enable rate-limiting on a public listener (throttle) inbound email for your enterprise. For more information on Default Host Access policies, see Sender Group Syntax.

The default host access policy for the listener is then accepted.

```
You are now going to configure how the appliance accepts mail by

creating a "Listener".

Please create a name for this listener (Ex: "InboundMail"):

[]> InboundMail



Please choose an IP interface for this Listener.

1. Management (192.168.42.42/24: mail3.example.com)

2. PrivateNet (192.168.1.1/24: mail3.example.com)

3. PublicNet (192.168.2.1/24: mail3.example.com)

[1]> 3



Enter the domains or specific addresses you want to accept mail for.
```

```
Hostnames such as "example.com" are allowed.

Partial hostnames such as ".example.com" are allowed.

Usernames such as "postmaster@" are allowed.

Full email addresses such as "joe@example.com" or "joe@[1.2.3.4]" are allowed.

Separate multiple addresses with commas.

[]> example.com


Would you like to configure SMTP routes for example.com?  [Y]> y


Enter the destination mail server which you want mail for example.com to be delivered.
Separate multiple entries with commas.

[]> exchange.example.com


Do you want to enable rate limiting for this listener? (Rate limiting defines the maximum number
of recipients per hour you are willing to receive from a remote domain.) [Y]> y


Enter the maximum number of recipients per hour to accept from a remote domain.

[]> 4500


Default Policy Parameters

===========================

Maximum Message Size: 100M

Maximum Number Of Connections From A Single IP: 1,000

Maximum Number Of Messages Per Connection: 1,000

Maximum Number Of Recipients Per Message: 1,000

Maximum Number Of Recipients Per Hour: 4,500

Maximum Recipients Per Hour SMTP Response:
```

```
      452 Too many recipients received this hour

Use SenderBase for Flow Control:  Yes

Virus Detection Enabled:  Yes

Allow TLS Connections: No

Would you like to change the default host access policy?  [N]> n


Listener InboundMail created.

Defaults have been set for a Public listener.

Use the listenerconfig->EDIT command to customize the listener.

*****
```

Private Listener

In this example portion of the systemsetup command, a private listener named OutboundMail is configured to run on the PrivateNet IP interface. Then, it is configured to relay all email for all hosts within the domain example.com . (Note the dot at the beginning of the entry: .example.com)

The default value for rate limiting (not enabled) and the default host access policy for this listener are then accepted.

Note that the default values for a private listener differ from the public listener created earlier. For more information, see Working with Listeners.

```
Do you want to configure the appliance to relay mail for internal hosts?  [Y]> y


Please create a name for this listener (Ex: "OutboundMail"):

[]> OutboundMail


Please choose an IP interface for this Listener.

1. Management (192.168.42.42/24: mail3.example.com)

2. PrivateNet (192.168.1.1/24: mail3.example.com)

3. PublicNet (192.168.2.1/24: mail3.example.com)

[1]> 2
```

```
Please specify the systems allowed to relay email through the appliance.

Hostnames such as "example.com" are allowed.

Partial hostnames such as ".example.com" are allowed.

IP addresses, IP address ranges, and partial IP addressed are allowed.

Separate multiple entries with commas.

[]> .example.com



Do you want to enable rate limiting for this listener?
(Rate limiting defines the maximum number of recipients per hour you are willing
to receive from a remote domain.)  [N]> n



Default Policy Parameters

==========================

Maximum Message Size: 100M

Maximum Number Of Connections From A Single IP: 600

Maximum Number Of Messages Per Connection: 10,000

Maximum Number Of Recipients Per Message: 100,000

Maximum Number Of Recipients Per Hour: Disabled

Use SenderBase for Flow Control:  No

Virus Detection Enabled:  Yes

Allow TLS Connections: No

Would you like to change the default host access policy?  [N]> n



Listener OutboundMAil created.

Defaults have been set for a Private listener.

Use the listenerconfig->EDIT command to customize the listener.

*****
```

Listener Example for C190 Appliances

✎

Note The following example of creating a listener applies to C170 and C190 appliances only.

In this example portion of the systemsetup command, a listener named MailInterface is configured to run on the MailNet IP interface. Then, it is configured to accept all email for the domain `example.com`. An initial SMTP route to the mail exchange `exchange.example.com` is configured. Then, the same listener is configured to relay all email for all hosts within the `domain example.com`. (Note the dot at the beginning of the entry: `.example.com`)

Rate limiting is enabled, and the maximum value of 450 recipients per hour from a single host is specified for the public listener.

✎

Note The value you enter for maximum recipients per hour you are willing to receive from a remote host is a completely arbitrary value, one that is usually relative to the size of the enterprise for which you are administering email. For example, a sender who sends 200 messages in an hour might be considered a "spammer" (sender of unsolicited bulk email), but if you are configuring the email gateway to handle all email for a 10,000 person company, 200 messages per hour from a remote host may be a reasonable value. Conversely, in a 50-person company, someone sending 200 messages in an hour to you may be an obvious spammer. You must choose an appropriate value when you enable rate-limiting on a public listener (throttle) inbound email for your enterprise. For more information on Default Host Access policies, see Sender Group Syntax.

The default host access policy for the listener is then accepted.

```
You are now going to configure how the appliance accepts mail by creating a "Listener".

Please create a name for this listener (Ex: "MailInterface"):

[]> MailInterface


Please choose an IP interface for this Listener.

1. MailNet (10.1.1.1/24: mail3.example.com)

2. Management (192.168.42.42/24: mail3.example.com)

[1]> 1


Enter the domain names or specific email addresses you want to accept mail for.


Hostnames such as "example.com" are allowed.
```

```
Partial hostnames such as ".example.com" are allowed.

Usernames such as "postmaster@" are allowed.

Full email addresses such as "joe@example.com" or "joe@[1.2.3.4]" are allowed.

Separate multiple addresses with commas.

[]> example.com


Would you like to configure SMTP routes for example.com?  [Y]> y


Enter the destination mail server where you want mail for example.com to be delivered.
Separate multiple entries with commas.

[]> exchange.example.com


Please specify the systems allowed to relay email through the appliance.

Hostnames such as "example.com" are allowed.

Partial hostnames such as ".example.com" are allowed.

IP addresses, IP address ranges, and partial IP addresses are allowed.

Separate multiple entries with commas.

[]> .example.com


Do you want to enable rate limiting for this listener?
(Rate limiting defines the maximum number of recipients per hour you are willing
to receive from a remote domain.)  [Y]> y


Enter the maximum number of recipients per hour to accept from a remote domain.

[]> 450


Default Policy Parameters

==========================
```

```
Maximum Message Size: 10M

Maximum Number Of Connections From A Single IP: 50

Maximum Number Of Messages Per Connection: 100

Maximum Number Of Recipients Per Message: 100

Maximum Number Of Recipients Per Hour: 450

Maximum Recipients Per Hour SMTP Response:

    452 Too many recipients received this hour

Use SenderBase for Flow Control:  Yes

Spam Detection Enabled:  Yes

Virus Detection Enabled:  Yes

Allow TLS Connections: No

Would you like to change the default host access policy?  [N]>



Listener MailInterface created.

Defaults have been set for a Public listener.

Use the listenerconfig->EDIT command to customize the listener.

*****
```

Note Because the systemsetup command only configures one listener for both inbound and outbound mail for C170 and C190 appliances, all outgoing mail will be calculated in the Mail Flow Monitor feature (which is normally used for inbound messages). See Using Email Security Monitor

Enable Anti-Spam

Your email gateway ships with a 30-day evaluation key for the Anti-Spam software. During this portion of the **systemsetup** command, you can choose to accept the license agreements and enable Anti-Spam globally on the email gateway.

Anti-Spam scanning will then be enabled on the incoming mail policy.

Note   If you do not accept the license agreement, Anti-Spam is not enabled on the email gateway.

See Managing Spam and Graymail for all of the Anti-Spam configuration options available on the email gateway.

Select a Default Anti-Spam Scanning Engine

If you have enabled more than one anti-spam scanning engine, you are prompted to select which engine will be enabled for use on the default incoming mail policy.

Enable the Spam Quarantine

If you choose to enable an anti-spam service, you can enable the incoming mail policy to send spam and suspected spam messages to the local Spam Quarantine. Enabling the Spam Quarantine also enables the end-user quarantine on the email gateway. Only administrators can access the end-user quarantine until end-user access is configured.

See Setting Up the Local Spam Quarantine .

Enable Anti-Virus Scanning

Your email gateway ships with a 30-day evaluation key for virus scanning engines. During this portion of the `systemsetup` command, you can choose to accept one or more license agreements and enable anti-virus scanning on the email gateway. You must accept a license agreement for each anti-virus scanning engine you want to enable on your email gateway.

After you accept the agreement, the anti-virus scanning engine you selected is enabled on the incoming mail policy. The email gateway scans incoming mail for viruses, but it does not repair infected attachments. The appliance drops infected messages.

See Anti-Virus for the anti-virus configuration options available on the email gateway.

Enable Outbreak Filters

This next step prompts you to enable Outbreak Filters. Your email gateway ships with a 30-day evaluation key for Outbreak Filters.

Related Topics

• Outbreak Filters, on page 37

Outbreak Filters

Outbreak Filters provide a "first line of defense" against new virus outbreaks by quarantining suspicious messages until traditional Anti-Virus security services can be updated with a new virus signature file. If enabled, Outbreak Filters will be enabled on the default Incoming Mail Policy.

If you choose to enable Outbreak Filters, enter a threshold value and whether you would like to receive Outbreak Filters alerts. For more information about Outbreak Filters and threshold values, see Outbreak Filters.

Configure the Alert Settings and AutoSupport

Cisco AsyncOS sends alert messages to a user via email if there is a system error that requires the user's intervention. Add at least one email address that receives system alerts. Separate multiple addresses with commas. The email addresses that you enter initially receive all types of alerts at all levels, except for Directory Harvest Attack Prevention alerts. You can add more granularity to the alert configuration later using the `alertconfig` command in the CLI or the **System Administration > Alerts** page in the GUI. For more information, see section *Alerts* of chapter *Distributing Administrative Tasks* of the *Cisco Secure Email Gateway Guide*.

The AutoSupport feature keeps the Cisco Customer Support team aware of issues with your email gateway so that Cisco can provide industry-leading support to you. Answer "Yes" to send Cisco support alerts and weekly status updates. (For more information, see section *AutoSupport* of chapter *Distributing Administrative Tasks* of the *Cisco Secure Email Gateway Guide*.)

Configure Scheduled Reporting

Enter an address to which to send the default scheduled reports. You can leave this value blank and the reports will be archived on the email gateway instead of sent via email.

Configure Time Settings

Cisco AsyncOS allows you to use the Network Time Protocol (NTP) to synchronize time with other servers on your network or the Internet, or to manually set the system clock. You must also set the time zone on the email gateway so that timestamps in message headers and log files are correct. You can also use the Cisco Systems time servers to synchronize the time on your email gateway.

Choose the Continent, Country, and Timezone and whether to use NTP including the name of the NTP server to use.

Commit Changes

Finally, the System Setup Wizard will ask you to `commit` the configuration changes you have made throughout the procedure. Answer "Yes" if you want to commit the changes.

When you have successfully completed the System Setup Wizard, the following message will appear and you will be presented with the command prompt:

```
Congratulations! System setup is complete. For advanced configuration, please refer to the
 User Guide.
```

```
mail3.example.com>
```

The email gateway is now ready to send email.

Test the Configuration

To test the Cisco AsyncOS configuration, you can use the mailconfig command immediately to send a test email containing the system configuration data you just created with the `systemsetup` command:

```
mail3.example.com> mailconfig
```

```
Please enter the email address to which you want to send
```

```
the configuration file. Separate multiple addresses with commas.

[]> user@example.com



The configuration file has been sent to user@example.com.

mail3.example.com>
```

Send the configuration to a mailbox to which you have access to confirm that the system is able to send email on your network.

Immediate Alerts

The email gateway uses feature keys to enable features. The first time you create a listener in the `systemsetup` command, enable Anti-Spam, enable Sophos or McAfee Anti-Virus, or enable Outbreak Filters, an alert is generated and sent to the addresses you specified in Step 2: System, on page 18.

The alert notifies you periodically of the time remaining on the key. For example:

```
Your "Receiving" key will expire in under 30 day(s).
Please contact IronPort Customer Support.


Your "Sophos" key will expire in under 30 day(s).
Please contact IronPort Customer Support.


Your "Outbreak Filters" key will expire in under 30 day(s).
Please contact IronPort Customer Support.
```

For information on enabling a feature beyond the 30-day evaluation period, contact your Cisco sales representative. You can see how much time remains on a key via the **System Administration > Feature Keys** page or by issuing the `featurekey` command. (For more information, see Feature Keys.)

Configuring your system as an Enterprise Gateway

To configure your system as an Enterprise Gateway (accepting email from the Internet), complete this chapter first, and then see Configuring the Gateway to Receive Email for more information.

Verifying Your Configuration and Next Steps

Now that system setup is complete, your email gateway should be sending and receiving email. If you have enabled the anti-virus, anti-spam, and virus-outbreak filters security features, the system will also be scanning incoming and outgoing mail for spam and viruses.

The next step is to understand how to customize your email gateway's configuration. Understanding the Email Pipeline provides a detailed overview of how email is routed through the system. Each feature is processed in order (from top to bottom) and is described in the remaining chapters of this guide.