Cisco IronPort AsyncOS 7.6 for Email Configuration Guide

February 6, 2012

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-26342-01

Cisco IronPort AsyncOS 7.6 for Email Configuration Guide

CONTENTS

CHAPTER 1

Getting Started with the Cisco IronPort Email Security Appliance

What's New in This Release

This section describes the new features and enhancements in AsyncOS for Email Security 7.6. For more information about the release, see the product release notes, which are available on the Cisco IronPort Customer Support page at the following URL:

`http://www.cisco.com/web/ironport/index.html`

You might also find it useful to review release notes for earlier releases to see the features and enhancements that were previously added. To view those release notes on the Support Portal, click the Earlier Releases link on the appropriate appliance documentation page.

New Feature: IPv6 Support

AsyncOS 7.6 adds Internet Protocol Version 6 (IPv6) address compatibility to your Email Security appliance. You can use both IPv4 and IPv6 addresses for your appliance's IP interfaces. IPv6 addresses are also an option for the following features:

- Gateways (default routers) and static routes.
- SMTP routes.
- SMTP Call Ahead.
- Trace.
- Senders for Host Access Tables.
- Recipients for Recipient Access Tables.
- Content Filter's Remote IP condition and Send to Alternate Destination Host action.
- Destination Controls, where you can specify whether IPv4 or IPv6 addresses are preferred.

- Outbreak Filters' Bypass Domain Scanning field.
- Report searches.

AsyncOS supports the following formats for IPv6 addresses:

- `2620:101:2004:4202::0-2620:101:2004:4202::ff`
- `2620:101:2004:4202::`
- `2620:101:2004:4202::23`
- `2620:101:2004:4202::/64`

New Feature: RSA Enterprise Manager Integration

AsyncOS 7.6's RSA Enterprise Manager Integration allows your organization to migrate an Email Security appliance's Data Loss Prevention policies to RSA Security's Enterprise Manager software in order to distribute those policies to all of your vectors enforcement. With RSA Enterprise Manager Integration, you can ensure consistent DLP policies across your enterprise and still have the option to manage policies on a local Email Security appliance when needed. For users of RSA's DLP Datacenter, RSA Enterprise Manager Integration also provides fingerprinting detection for scanning source code and documents to certain DLP policies.

Enterprise Manager is a third-party software offered by RSA Security, Inc. It is not a part of the Cisco IronPort Email Security appliance.

See the Chapter 11, "Data Loss Prevention" for more information.

As part of RSA Enterprise Manager Integration, AsyncOS now includes a User Distinguished Name LDAP query for LDAP profiles. This query returns a message sender's distinguished name for the Email Security appliance to include with all the other DLP incident data it sends to Enterprise Manager. See the "LDAP Queries" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information.

Enhancement: DLP Message Actions

Starting in AsyncOS 7.6, the primary and secondary actions performed by DLP policies are now defined as *message actions*. You create message actions using the Mail Policies > DLP Message Actions page in the GUI and then add the actions to your DLP policies. When updating from a previous version of AsyncOS, the system automatically generates new message actions based on the primary and secondary actions defined in your existing DLP policies.

See the Chapter 11, "Data Loss Prevention" for more information.

Enhancement: DLP Message Tracking Privileges By User Group

AsyncOS 7.6 allows you to choose which non-administrator user can view sensitive DLP-related information in Message Tracking by user role. See the "Common Administrative Tasks" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide* for more information.

Enhancement: RSA Email DLP's "Quarantine a Copy and Deliver" Option

AsyncOS 7.6 provides an option to quarantine a copy of a message that violates a RSA Email DLP policy while still delivering the original message.

See DLP Policies, page 11-10 for more information.

Enhancement: SenderBase Reputation Service Requires an Anti-Spam Feature Key

Starting in AsyncOS 7.6, an Email Security appliance requires an anti-spam system feature key in order to use the SenderBase Reputation Service.

New Feature: DKIM Verification Profiles

AsyncOS 7.6 adds DKIM verification profiles, which are lists of parameters that the Email Security appliance's mail flow policies use for verifying DKIM signatures. For example, you can create two verification profiles, one that allows 30 seconds before a query times out and a second that allows only 3 seconds before a query times out. You can assign the second verification profile to the Throttled mail flow policy to prevent connection starvation in case of a DDoS.

See the "Email Authentication" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information.

Enhancement: New Tags for DKIM Signing Profiles

AsyncOS 7.6 adds a new list of tags to include in DKIM message signatures. You select which tags you want to include in the signatures when creating a DKIM signing profile. The following tags are available:

- **"i" Tag.** The identity of the user or agent (e.g., a mailing list manager) on whose behalf the message is signed.
- **"q" Tag.** A comma-separated list of query methods used to retrieve the public key.
- **"t" Tag.** The timestamp of when the signature was created.
- **"x" Tag.** The expiration time of the signature, in seconds. (The option in include "x" tag information existed in previous versions of AsyncOS 7.6.)
- **"z" Tag.** A vertical bar-separated (i.e., |) list of header fields present when the message was signed.

See the "Email Authentication" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information.

New Feature: DKIM Signing of System-Generated Messages

AsyncOS 7.6 allows you to choose whether to sign system-generated messages with a DKIM signature. The types of system-generated message that the Email Security appliance will sign include the following:

- Cisco IronPort Spam Quarantine notifications
- Content filter-generated notifications

- Configuration messages
- Support requests

See the "Email Authentication" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information.

Enhancement: Skip DKIM Signing Action

In AsyncOS 7.6, content filters now include an action to skip DKIM signing.

See Content Filter Actions, page 6-12 for more information.

Enhancement: Rate Limiting and Enforced TLS for Envelope Senders in Mail Flow Policies

AsyncOS 7.6 updates Mail Flow Policies with the option to limit number of recipients during a specified time period that a listener will receive from a unique envelope sender, based on the mail-from address. Each listener tracks its own rate limiting threshold; however, because all listeners validate against a single counter, it is more likely that the rate limit will be exceeded if messages from the same mail-from address are received by multiple listeners.

You can also make TLS connections mandatory for envelope senders from a certain domain or with a specific email address when the mail flow policy has a setting of Preferred for encryption over TLS.

See Mail Flow Policies: Access Rules and Parameters, page 5-8 for more information.

You specify the domains and email addresses for these enevelop senders using an address list. See Address Lists, page 5-39 for more information.

AsyncOS also adds a Rate Limiting report that allows you to quickly identify individual senders of large numbers of messages. Use this report to help you to control spam from internal user accounts, identify compromised user accounts, limit out-of-control applications that use email, and avoid damaging your organization's online reputation and the attendant hassles resulting from this situation.

See the "Using Email Security Monitor" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide* for more information.

Enhancement: Separate Update Servers for AsyncOS Upgrades and Other Service Updates

AsyncOS 7.6 allows you to specify a different update server for AsyncOS upgrades than the one used for other service updates, such as feature key updates, outbreak filters, and time zone rules. For example, you can specify a local server for downloading AsyncOS upgrades while using the Cisco IronPort update servers for the other service updates.

See Service Updates, page 15-10 for more information.

Enhanced: Web User Interface Protection

AsyncOS 7.6 for Email includes additional protection from cross-site request forgeries (CSRF) and other attacks on the web user interface.

The Email Security Appliance Documentation Set

The documentation for the Email Security appliance includes the following books:

- *Cisco IronPort AsyncOS for Email Daily Management Guide*. This guide provides instructions for performing common, everyday tasks that system administrators use to manage and monitor the Cisco IronPort appliance, such as viewing email traffic using the Email Security Monitor, tracking email messages, managing system quarantines, and troubleshooting the appliance. It also provides reference information for features that system administrators interact with on a regular basis, including Email Security Monitor pages, AsyncOS logs, CLI support commands, and quarantines.
- *Cisco IronPort AsyncOS for Email Configuration Guide*. This guide is recommended for system administrators who are setting up a new Cisco IronPort appliance and want to learn about its email delivery features. It provides instructions on installing the appliance into an existing network infrastructure and setting it up as an email gateway appliance. It also includes reference information and configuration instructions for email delivery features such as the Email Pipeline, Outbreak Filters, content filters, RSA Email DLP, email encryption, anti-virus scanning, and anti-spam scanning.
- *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*. This guide provides instructions configuring the advanced features of the Cisco IronPort appliance. Topics include configuring the appliance to work with LDAP, creating message filters to enforce email policies, organizing multiple appliances into clusters, and customizing the listeners on the appliance. In addition to configuration, this guide provides reference material for advanced features such as message filter rules and actions, regular expressions used in content dictionaries and message filter rules, and LDAP query syntax and attributes.
- *Cisco IronPort AsyncOS CLI Reference Guide.* This guide provides a detailed list of the commands in the AsyncOS command line interface (CLI), as well as examples of the commands in use. System administrators can use this guide for reference when using the CLI on the Cisco IronPort appliance.

Occasionally, this book refers to the other guides for additional information about topics. These guides are available on the Documentation CD that came with your Cisco IronPort appliance as well as the Cisco IronPort Customer Support Portal. For more information, see Cisco IronPort Support Community, page 1-10.

How to Use This Guide

Use this guide as a resource to learn about the features of your Cisco IronPort appliance. The topics are organized in a logical order. You might not need to read every chapter in the book. Review the Table of Contents and the section called How This Book Is Organized, page 1-6 to determine which chapters are relevant to your system.

You can also use this guide as a reference book. It contains important information, such as network and firewall configuration settings, that you can refer to throughout the life of the appliance.

The guide is distributed in print and electronically as PDF and HTML files. The electronic versions of the guide are available on the Cisco IronPort Customer Support Portal. You can also access the HTML online help version of the book in the appliance GUI by clicking the Help and Support link in the upper-right corner.

Before You Begin

Before you read this guide, review the *Cisco IronPort Quickstart Guide* and the latest product release notes for your appliance. In this guide, it is assumed that you have unpacked the appliance, physically installed it in a rack, and turned it on.

Note If you have already cabled your appliance to your network, ensure that the default IP address for the Cisco IronPort appliance does not conflict with other IP addresses on your network. The IP address that is pre-configured on the Management port (on Cisco IronPort X1000/1000T/1050/1060/1070, C60/600/650/660/670, and C30/300/300D/350/350D/360/370 appliances) or the Data 1 port (on Cisco IronPort C10/100/150/160 appliances) is `192.168.42.42`.

How This Book Is Organized

Chapter 1, "Getting Started with the Cisco IronPort Email Security Appliance" provides an introduction to the Cisco IronPort appliance and defines its key features and role in the enterprise network. New features of the current release are described.

Chapter 2, "Overview" introduces Cisco IronPort AsyncOS for Email and discusses administration of the Cisco IronPort appliance through its GUI and CLI. Conventions for using the CLI are described. This chapter also contains an overview of general purpose CLI commands.

Chapter 3, "Setup and Installation" describes the options for connecting to the Cisco IronPort appliance, including network planning, and initial system setup and configuration of the appliance.

Chapter 4, "Understanding the Email Pipeline" provides an overview of the email pipeline — the flow that email follows as it is processed by the Cisco IronPort appliance — and brief descriptions of the features that comprise the pipeline. The descriptions include cross-references to the sections containing detailed explanations of the features.

Chapter 5, "Configuring the Gateway to Receive Email" describes the process of configuring the appliance as an email gateway. This chapter introduces the concepts of interfaces, listeners, and the Host Access Table (HAT) — which support incoming email traffic and the Mail Flow Monitor.

Chapter 6, "Email Security Manager" describes Email Security Manager, the single, comprehensive dashboard to manage all email security services and applications on Cisco IronPort appliances. Email Security Manager allows you to manage the Outbreak Filters feature, Anti-Spam, Anti-Virus, and email content policies — on a per-recipient or per-sender basis, through distinct inbound and outbound policies.

Chapter 7, "Reputation Filtering" provides an overview of how SenderBase Reputation Service scores are used to control incoming mail based on the reputation of the message sender.

Chapter 9, "Anti-Spam" describes the unique approach to fighting spam with the SenderBase Reputation Filters, Cisco IronPort Anti-Spam, and Cisco IronPort Intelligent Multi-Scan features integrated into the Cisco IronPort appliance.

Chapter 8, "Anti-Virus" explains the Sophos and McAfee Anti-Virus scanning features integrated into the Cisco IronPort appliance.

Chapter 10, "Outbreak Filters" explains how Outbreak Filters proactively provide a critical first layer of defense against new virus, scam, and phishing outbreaks. By detecting new outbreaks in real-time and dynamically responding to prevent suspicious traffic from entering the network, Outbreak Filters offer protection until new signature updates are deployed.

Chapter 11, "Data Loss Prevention" describes how to use the data loss prevention features from RSA Security, Inc. to protect your organization's information and intellectual property, as well as enforce regulatory and organizational compliance by preventing users from unintentionally emailing sensitive data.

Chapter 12, "Cisco IronPortEmail Encryption" describes the process you use to encrypt email using the Cisco IronPort Encryption appliance or the hosted key service.

Chapter 13, "SenderBase Network Participation" describes how to share data from your appliance with the SenderBase Network.

Chapter 14, "Text Resources" details creating text resources such as content dictionaries, notification templates, and disclaimers for use in various components of AsyncOS.

Chapter 15, "System Administration" describes typical administration commands for managing and monitoring the Cisco IronPort appliance, such as working with feature keys, upgrading AsyncOS, reverting AsyncOS, and performing routine system maintenance. Maintenance tasks include setting the system time, changing the administrator password, and taking the system offline. This chapter also describes how to configure the network operation of the Cisco IronPort appliance, including DNS, interface, routing, and hostname settings.

Chapter 16, "Enabling Your C350D Appliance" describes the Cisco IronPort C300D, C350D, and C360D appliances.

Chapter 17, "The Cisco IronPort M-Series Security Management Appliance" describes the Cisco IronPort M-Series appliance, which is designed to centralize and consolidate important policy and runtime data, providing administrators and end users with a single interface for managing reporting and auditing information.

Appendix A, "Accessing the Appliance" describes how to access the Cisco IronPort appliance for uploading and downloading files.

Appendix B, "Assigning Network and IP Addresses" describes general rules on networks and IP address assignments and presents strategies for connecting the Cisco IronPort appliance within an enterprise network infrastructure.

Appendix C, "Firewall Information" describes the possible ports that may need to be opened for proper operation of the Cisco IronPort appliance behind a security firewall.

Appendix D, "Cisco IronPort Systems, LLC Software License Agreement" includes the software license agreement for the Cisco IronPort Email Security appliance.

Topics Discussed in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*

The following topics are discussed in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*:

Chapter 1, "Customizing Listeners" describes the process for tailoring the configuration of your Enterprise Email Gateway. This chapter discusses, in detail, advanced features available to you as you configure interfaces and listeners to handle email receiving through the gateway.

Chapter 2, "Configuring Routing and Delivery Features" explains the features that affect email routing and delivery of email traveling through the Cisco IronPort appliance.

Chapter 3, "LDAP Queries" describes how your Cisco IronPort appliance can connect to your corporate Lightweight Directory Access Protocol (LDAP) servers and perform queries for the purposes of verifying recipients to accept (including group membership), mail routing and address rewriting. masquerading headers, and supporting for SMTP authentication.

Chapter 4, “Email Authentication” details the process of configuring and enabling email authentication on an Cisco IronPort appliance. Cisco IronPort AsyncOS supports several types of email authentication, including Sender Policy Framework (SPF), Sender ID Framework (SIDF), and DomainKeys Identified Mail (DKIM) verification of incoming mail, as well as DomainKeys and DKIM signing of outgoing mail.

Chapter 5, “Using Message Filters to Enforce Email Policies” describes how to use Message Filters to define rules for handling email, including the ability to modify the content of messages through the attachment filtering, image analysis, and content dictionary features.

Chapter 7, “Advanced Network Configuration” includes information about NIC pairing, virtual LANs and more.

Chapter 8, “Centralized Management” describes the centralized management feature, which allows you to manage and configure multiple appliances. The centralized management feature provides increased reliability, flexibility, and scalability within your network, allowing you to manage globally while complying with local policies.

Appendix A, “AsyncOS Quick Reference Guide” provides a quick reference for most commands in the CLI.

Appendix B, “Accessing the Appliance” describes how to access the Cisco IronPort appliance to send and retrieve files from Cisco IronPort appliance.

The following topics are discussed in the *Cisco IronPort AsyncOS for Email Daily Management Guide*

Chapter 1, “Managing the Cisco IronPort Email Appliance,” provides an introduction to the Cisco IronPort appliance and defines its key features and role in the enterprise network.

Chapter 2, “Using Email Security Monitor,” describes the Mail Flow Monitor feature: a powerful, web-based console that provides complete visibility into all inbound email traffic for your enterprise.

Chapter 3, “Tracking Email Messages,” describes local message tracking. You can use message tracking to determine if a particular message was delivered, found to contain a virus, or placed in a spam quarantine.

Chapter 4, “Quarantines,” describes the special queues or repositories used to hold and process messages. Messages in quarantines can be delivered or deleted, based on how you configured the quarantine. This includes the Cisco IronPort Spam quarantine.

Chapter 5, “Logging,” describes the logging and log subscription functionality of the Cisco IronPort appliance.

Chapter 6, “Managing and Monitoring via the CLI,” describes the commands available in the CLI available to you as you monitor the mail flow through the gateway.

Chapter 7, “Other Tasks in the GUI,” describes typical administration tasks for managing and monitoring the Cisco IronPort appliance through the GUI.

Chapter 8, “Common Administrative Tasks,” describes typical administration commands for managing and monitoring the Cisco IronPort appliance, such adding users, managing the configuration file, and managing SSH keys. This chapter also describes how to request technical support, allow Cisco IronPort customer support remote access to your appliance, and use feature keys.

Chapter 9, “Testing and Troubleshooting” describes the process of creating so-called *black hole listeners* for testing the system performance and troubleshooting configuration problems.

Appendix A, "Accessing the Appliance," describes how to access the Cisco IronPort appliance for uploading and downloading files.

Typographic Conventions

Typeface	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output.	`Please choose an IP interface for this Listener.` The `sethostname` command sets the name of the Cisco IronPort appliance.
AaBbCc123	User input, in contrast to on-screen computer output.	`mail3.example.com>` **`commit`** `Please enter some comments describing your changes:` `[]>` **`Changed the system hostname`**
AaBbCc123	Book titles, new terms, emphasized words, and command line variables; for command line variables, the italicized text is a placeholder for the actual name or value.	Read the *Cisco IronPort Quickstart Guide*. The Cisco IronPort appliance *must* be able to uniquely select an interface to send an outgoing packet. `Before you begin, please reset your password to a new value.` `Old password:` **`ironport`** `New password:` *`your_new_password`* `Retype new password:` **`your_new_password`**

Where to Find More Information

Cisco offers the following resources to learn more about the Cisco IronPort Email Security appliance.

Cisco IronPort Technical Training

Cisco IronPort Systems Technical Training Services can help you acquire the knowledge and skills necessary to successfully evaluate, integrate, deploy, maintain, and support Cisco IronPort security products and solutions.

Use one of the following methods to contact Cisco IronPort Technical Training Services:

Training. For question relating to registration and general training:

- http://training.ironport.com
- training@ironport.com

Certifications. For questions relating to certificates and certification exams:

- http://training.ironport.com/certification.html
- certification@ironport.com

Knowledge Base

You can access the Cisco IronPort Knowledge Base on the Customer Support Portal at the following URL:

`http://www.cisco.com/web/ironport/knowledgebase.html`

Note You need a Cisco.com User ID to access the site. If you do not have a Cisco.com User ID, you can register for one here: `https://tools.cisco.com/RPF/register/register.do`

The Knowledge Base contains a wealth of information on topics related to Cisco IronPort products.

Articles generally fall into one of the following categories:

- **How-To.** These articles explain how to do something with a Cisco IronPort product. For example, a how-to article might explain the procedures for backing up and restoring a database for an appliance.
- **Problem-and-Solution.** A problem-and-solution article addresses a particular error or issue that you might encounter when using a Cisco IronPort product. For example, a problem-and-solution article might explain what to do if a specific error message is displayed when you upgrade to a new version of the product.
- **Reference.** Reference articles typically provide lists of information, such as the error codes associated with a particular piece of hardware.
- **Troubleshooting.** Troubleshooting articles explain how to analyze and resolve common issues related to Cisco IronPort products. For example, a troubleshooting article might provide steps to follow if you are having problems with DNS.

Each article in the Knowledge Base has a unique answer ID number.

Cisco IronPort Support Community

The Cisco IronPort Support Community is an online forum for Cisco IronPort customers, partners, and employees. It provides a place to discuss general email and web security issues, as well as technical information about specific Cisco IronPort products. You can post topics to the forum to ask questions and share information with other Cisco IronPort users.

You access the Cisco IronPort Support Community on the Customer Support Portal at the following URL:

`https://supportforums.cisco.com`

Cisco IronPort Customer Support

You can request Cisco IronPort product support by phone, email, or online 24 hours a day, 7 days a week.

During customer support hours (24 hours per day, Monday through Friday excluding U.S. holidays), an engineer will contact you within an hour of your request.

To report a critical issue that requires urgent assistance outside of our office hours, please contact Cisco IronPort using one of the following methods:

U.S. toll-free: 1(877) 641- 4766

International: `http://cisco.com/web/ironport/contacts.html`

Support Site: `http://cisco.com/web/ironport/index.html`

If you purchased support through a reseller or another supplier, please contact that supplier directly with your product support issues.

Third Party Contributors

Some software included within Cisco IronPort AsyncOS is distributed under the terms, notices, and conditions of software license agreements of FreeBSD, Inc., Stichting Mathematisch Centrum, Corporation for National Research Initiatives, Inc., and other third party contributors, and all such terms and conditions are incorporated in Cisco IronPort license agreements.

The full text of these agreements can be found here:

https://support.ironport.com/3rdparty/AsyncOS_User_Guide-1-1.html.

Portions of the software within Cisco IronPort AsyncOS is based upon the RRDtool with the express written consent of Tobi Oetiker.

Portions of this document are reproduced with permission of Dell Computer Corporation. Portions of this document are reproduced with permission of McAfee, Inc. Portions of this document are reproduced with permission of Sophos Plc.

Cisco IronPort Welcomes Your Comments

The Cisco IronPort Technical Publications team is interested in improving the product documentation. Your comments and suggestions are always welcome. You can send comments to the following email address:

`contentsecuritydocs@cisco.com`

Please include the following part number in the subject of your message: OL-26342-01.

Cisco IronPort Email Security Appliance Overview

The Cisco IronPort Email Security appliance is a high-performance appliance designed to meet the email infrastructure needs of the most demanding enterprise networks. The Email Security appliance eliminates spam and viruses, enforces corporate policy, secures the network perimeter, and reduces the total cost of ownership (TCO) of enterprise email infrastructure.

Cisco IronPort Systems combines hardware, a hardened operating system, application, and supporting services to produce a purpose-built, rack-mount server appliance dedicated for enterprise messaging.

The Cisco IronPort AsyncOS™ operating system integrates several intelligent features into the Cisco IronPort appliance:

- **Anti-Spam** at the gateway, through the unique, multi-layer approach of SenderBase Reputation Filters and Cisco IronPort Anti-Spam integration.
- **Anti-Virus** at the gateway with the Sophos and McAfee Anti-Virus scanning engines.
- **Outbreak Filters**™, Cisco IronPort's unique, preventive protection against new virus, scam, and phishing outbreaks that can quarantine dangerous messages until new updates are applied, reducing the window of vulnerability to new message threats.
- **Spam Quarantine** either on-box or off, providing end user access to quarantined spam and suspected spam.

- **Email Authentication**. Cisco IronPort AsyncOS supports various forms of email authentication, including Sender Policy Framework (SPF), Sender ID Framework (SIDF), and DomainKeys Identified Mail (DKIM) verification of incoming mail, as well as DomainKeys and DKIM signing of outgoing mail.
- **Cisco IronPort Email Encryption**. You can encrypt outgoing mail to address HIPAA, GLBA and similar regulatory mandates. To do this, you configure an encryption policy on the Email Security appliance and use a local key server or hosted key service to encrypt the message.
- **Email Security Manager**, a single, comprehensive dashboard to manage all email security services and applications on the appliance. Email Security Manager can enforce email security based on user groups, allowing you to manage Cisco IronPort Reputation Filters, Outbreak Filters, Anti-Spam, Anti-Virus, and email content policies through distinct inbound and outbound policies.
- **On-box Quarantine areas** to hold messages that violate email policies. Quarantines seamlessly interact with the Outbreak Filters feature.
- **On-box message tracking**. AsyncOS for Email includes an on-box message tracking feature that makes it easy to find the status of messages that the Email Security appliance processes.
- **Mail Flow Monitoring** of all inbound and outbound email that provides complete visibility into all email traffic for your enterprise.
- **Access control** for inbound senders, based upon the sender's IP address, IP address range, or domain.
- Extensive **message filtering** technology allows you to enforce corporate policy and act on specific messages as they enter or leave your corporate infrastructure. Filter rules identify messages based on message or attachment content, information about the network, message envelope, message headers, or message body. Filter actions allow messages to be dropped, bounced, archived, blind carbon copied, or altered, or to generate notifications.
- **Message encryption via secure SMTP over Transport Layer Security** ensures messages travelling between your corporate infrastructure and other trusted hosts are encrypted.
- **Virtual Gateway**™ technology allows the Cisco IronPort appliance to function as several email gateways within a single server, which allows you to partition email from different sources or campaigns to be sent over separate IP addresses. This ensures that deliverability issues affecting one IP address do not impact others.

AsyncOS for Email is a proprietary operating system that has been highly optimized for the task of Internet messaging. AsyncOS is a "hardened" operating system: all unnecessary services have been removed, which increases security and optimizes system performance. Cisco IronPort stackless threading technology eliminates allocation of a dedicated memory stack to each task, which increases concurrency and stability of the MTA. The custom I/O-driven scheduler is optimized for massively concurrent I/O events required by the email gateway versus the preemptive time slicing of the CPU in traditional operating systems. AsyncFS, the file system underlying AsyncOS, is optimized for millions of small files and ensures data recoverability in the case of system failure.

AsyncOS for email supports RFC 2821-compliant Simple Mail Transfer Protocol (SMTP) to accept and deliver messages. The Cisco IronPort appliance is designed to be easy to configure and manage. Most reporting, monitoring, and configuration commands are available through both the web-based GUI via HTTP or HTTPS. In addition, an interactive Command Line Interface (CLI) which you access from a Secure Shell (SSH), telnet, or direct serial connection is provided for the system. The Cisco IronPort appliance also features a robust logging capability, allowing you to configure log subscriptions spanning the functionality of the entire system and reducing the time spent finding the information you need.

Mail Flow and the Cisco IronPort M-Series Appliance

If you include an M-Series appliance in your configuration, mail is sent to the Cisco IronPort M-Series appliance from other Cisco IronPort (C- and X-Series) appliances. A Cisco IronPort appliance that is configured to send mail to a Cisco IronPort M-Series appliance will automatically expect to receive mail released from the M-Series appliance and will not re-process those messages when they are received back — messages will bypass the HAT and other policy or scanning settings and be delivered. For this to work, the IP address of the Cisco IronPort M-Series appliance must not change. If the IP address of the Cisco IronPort M-Series appliance changes, the receiving C- or X-Series appliance will process the message as it would any other incoming message. Always use the same IP address for receiving and delivery on the Cisco IronPort M-Series appliance.

The Cisco IronPort M-Series appliance accepts mail for quarantining from the IP addresses specified in the Cisco IronPort Spam Quarantine settings. To configure the local quarantine on the Cisco IronPort M-Series appliance see the *Cisco IronPort AsyncOS for Security Management User Guide* . Note that the local quarantine on the Cisco IronPort M-Series appliance is referred to as an *external* quarantine by the other Cisco IronPort appliances sending mail to it.

Mail released by the Cisco IronPort M-Series appliance is delivered to the primary and secondary hosts (Cisco IronPort appliance or other groupware host) as defined in the Spam Quarantine Settings (see the *Cisco IronPort AsyncOS for Security Management User Guide*). Therefore, regardless of the number of Cisco IronPort appliances delivering mail to the Cisco IronPort M-Series appliance, all released mail, notifications, and alerts are sent to a single host (groupware or Cisco IronPort appliance). Take care to not overburden the primary host for delivery from the Cisco IronPort M-Series appliance.

CHAPTER 2

Overview

This chapter introduces the Cisco IronPort AsyncOS operating system and administration of the Cisco IronPort appliance through both the web-based Graphical User Interface (GUI) and Command Line Interface (CLI). Conventions for using each interface are described. This chapter also contains general-purpose CLI commands.

Web-based Graphical User Interface (GUI)

The graphical user interface (GUI) is the web-based alternative to the command line interface (CLI) for system monitoring and configuration. The GUI enables you to monitor the system using a simple web-based interface without having to learn the Cisco IronPort AsyncOS command syntax.

The GUI contains most of the functionality you need to configure and monitor the system. However, not all CLI commands are available in the GUI; some features are *only* available through the CLI. Many of the tasks listed throughout this book demonstrate how to accomplish a task from the GUI (when possible) first, followed by the CLI commands to accomplish the same task.

In the following chapters, you will learn how to use the GUI to:

- access the System Setup Wizard to perform the initial installation and configuration of the Cisco IronPort appliance.
- access Email Security Manager to enforce email security based on user groups, allowing you to manage Cisco IronPort Reputation Filters, Outbreak Filters, Anti-Spam, Anti-Virus, and email content filtering policies through distinct inbound and outbound policies.
- edit the Host Access Table (HAT) for a listener, customizing your own sender groups (updating whitelists, blacklists, and greylists) and tailoring mail flow policies by querying for a sender's reputation, including the SenderBase Reputation Score (SBRS).
- create and manage dictionaries, disclaimers, and other text resources.
- configure an encryption profile to use Cisco IronPort Email Encryption to encrypt outbound emails.
- configure global settings for Cisco IronPort Anti-Spam, Sophos Anti-Virus, Outbreak Filters, and SenderBase Network Participation.
- view status through XML pages, or access XML status information programmatically.

Browser Requirements

To access the web-based UI, your browser must support and be enabled to accept JavaScript and cookies, and it must be able to render HTML pages containing Cascading Style Sheets (CSS).

Note Beginning with AsyncOS 5.5, the web-based UI incorporates libraries from the Yahoo! User Interface (YUI) Library, which is a set of utilities and controls, written in JavaScript, for building richly interactive web applications. The purpose of this change is to provide an improved user experience in the web-based UI.

The YUI library supports the vast majority of browsers that are in general use. The YUI library also has a comprehensive, public approach to browser support and is committed to making sure that components work well in all of what are designated as "A-Grade" browsers. For more information on graded browser support, see:

`http://developer.yahoo.com/yui/articles/gbs/`

Cisco IronPort tests our web application with and recommends the following list of A-grade browsers to access the web-based UI:

- Firefox 3.6
- Windows XP and Vista: Internet Explorer 7 and 8
- Windows 7: Internet Explorer 8 and 9, Google Chrome, Firefox 4
- Mac OS X: Safari 4 and later, Firefox 4

Please note that when accessing the GUI, do not use multiple browser windows or tabs simultaneously to make changes to the Cisco IronPort appliance. Do not use concurrent GUI and CLI sessions either. Doing so will cause unexpected behavior and is not supported.

You may need to configure your browser's pop-up blocking settings in order to use the GUI because some buttons or links in the interface will cause additional windows to open.

Accessing the GUI

By default, the system ships with HTTP enabled on the Management interface (for Cisco IronPort C60/600/650/660/670, C30/300/350/360/370, and X1000/1050/1060/1070 appliances) or Data 1 (Cisco IronPort C10/100/150/160) interface. (For more information, see Enabling the GUI on an Interface, page -442.)

To access the GUI on a brand new system, access the following URL:

`http://192.168.42.42`

When the login page is displayed, log in to the system using the default username and password:

Factory Default Username and Password

- Username: **`admin`**
- Password: **`ironport`**

For example:

Figure 2-1 The Login Screen

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

During the initial system setup, you choose IP addresses for interfaces and whether to run HTTP and/or HTTPS services for those interfaces. When HTTP and/or HTTPS services have been enabled for an interface, you can use any supporting browser to view the GUI by entering the IP address or hostname of the IP interface as a URL in the location field ("address bar") of the browser. For example:

`http://192.168.1.1` or

`https://192.168.1.1` or

`http://mail3.example.com` or

`https://mail3.example.com`

Note If HTTPS has been enabled for an interface (and HTTP requests are *not* being redirected to the secure service), remember to access the GUI using the "`https://`" prefix.

Logging In

All users accessing the GUI must log in. Type your username and password, and then click Login to access the GUI. You must use a supported web browser (see Browser Requirements, page 2-2). You can log in with the admin account or with a specific user account you have created. (For more information, see "Adding Users" in the "Common Administrative Tasks" chapter of the *Cisco IronPort AsyncOS for Email Daily Management Guide*.)

After you have logged in, the Monitor > Incoming Mail Overview page is displayed.

GUI Sections and Basic Navigation

The GUI consists of the following menus which correspond to functions in your Cisco IronPort appliance: Monitor, Mail Policies, Security Services, Network, and System Administration. The following chapters will describe each section, including the tasks you perform on pages within each section.

Note Online help for the GUI is available from every page within the GUI. Click the **Help > Online Help** link at the top right of the page to access the online help.

You navigate among sections of the interface by clicking the menu headings for each main section (Monitor, Mail Policies, Security Services, Network, and System Administration). Within each menu are sub-sections that further group information and activities. For example, the Security Services section

contains the Anti-Spam section that lists the Anti-Spam pages. Accordingly, when referring to specific pages in the GUI, the documentation uses the menu name, followed by an arrow and then the page name. For example, **Security Services > SenderBase**.

Monitor menu

The Monitor section contain pages for the Mail Flow Monitor feature (Overview, Incoming Mail, Outgoing Destinations, Outgoing Senders, Delivery Status, Internal Users, Content Filters, Virus Outbreaks, Virus Types, System Capacity, System Status), Local and External Quarantines, and Scheduled Reports features. You can also access message tracking from this menu.

Mail Policies menu

The Mail Policies section contains pages for the Email Security Manager feature (including Mail Policies and Content Filters), the Host Access Table (HAT) and Recipient Access Table (RAT) configuration, Destination Controls, Bounce Verification, Domain Keys, Text Resources, and Dictionaries.

Security Services menu

The Security Services section contains pages to set global settings for the Anti-Spam, Anti-Virus, Cisco IronPort Email Encryption, Outbreak Filters, and SenderBase Network Participation features. You also enable the following features from this menu: Reporting, Message Tracking, External Spam Quarantine.

Network menu

The Network section contains pages for creating and managing IP interfaces, Listeners, SMTP Routes, DNS, Routing, Bounce Profiles, SMTP Authentication, and Incoming Relays.

System Administration menu

The System Administration section contains pages for the Trace, Alerting, User Management, LDAP, Log Subscription, Return Addresses, System Time, Configuration File management, Feature Key Settings, Feature Keys, Shutdown/Reboot, Upgrades, and System Setup Wizard features.

Centralized Management

If you have the Centralized Management feature and have enabled a cluster, you can browse machines in the cluster, create, delete, copy, and move settings among clusters, groups, and machines (that is, perform the equivalent of the `clustermode` and `clusterset` commands) from within the GUI.

For more information, see "Administering a Cluster from the GUI" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

The Commit Changes Button

The commit model in the GUI matches the same "explicit commit" model as used in the CLI. (For more information, see Committing Configuration Changes, page 2-9.) As you make configuration changes in the GUI, you now must explicitly commit those changes by clicking the **Commit Changes** button. This button displays when you have uncommitted changes that need to be saved.

Figure 2-2 The Commit Changes Button

Clicking the **Commit Changes** button displays a page where you can add a comment and commit the changes, abandon all changes made since the most recent commit (the equivalent of the `clear` command in the CLI; see Clearing Configuration Changes, page 2-10), or cancel.

Figure 2-3 Confirming Committed Changes

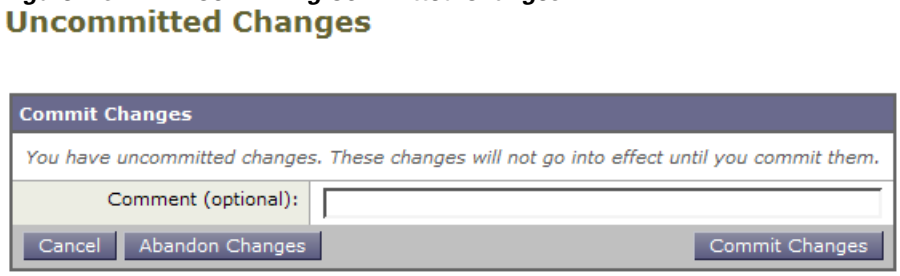

Viewing Active Sessions

From the GUI, you can view all users currently logged into the Email Security appliance and information about their sessions.

To view these active sessions, click **Options > Active Sessions** at the top right of the page.

From the Active Sessions page you can view the user name, the user role, the time the user logged in, idle time, and whether the user is logged in from the command line or the GUI.

Figure 2-4 Active Sessions

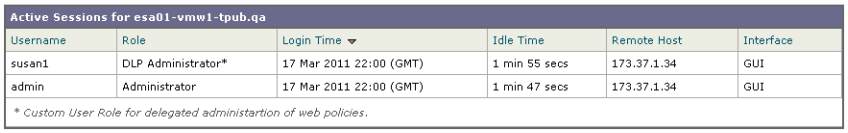

Command Line Interface (CLI)

The Cisco IronPort AsyncOS Command Line Interface is an interactive interface designed to allow you to configure and monitor the Cisco IronPort appliance. The commands are invoked by entering the command name with or without any arguments. If you enter the command without arguments, the command prompts you for the required information.

The Command Line Interface is accessible via SSH or Telnet on IP interfaces that have been configured with these services enabled, or via terminal emulation software on the serial port. By factory default, SSH and Telnet are configured on the Management port. Use the `interfaceconfig` command described in Configuring the Gateway to Receive Email, page 5-1 to disable these services.

For more information about specific CLI commands, see the *Cisco IronPort AsyncOS CLI Reference Guide*.

Command Line Interface Conventions

This section describes the rules and conventions of the AsyncOS CLI.

Command Prompt

The top-level command prompt consists of the fully qualified hostname, followed by the greater than (>) symbol, followed by a space. For example:

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

If the appliance has been configured as part of a cluster with the Centralized Management feature, the prompt in the CLI changes to indicate the current mode. For example:

```
(Cluster Americas) >
```

or

```
(Machine losangeles.example.com) >
```

See "Centralized Management" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information.

When running commands, the CLI requires input from you. When the CLI is expecting input from you, the command prompt shows the default input enclosed in square brackets ([]) followed by the greater than (>) symbol. When there is no default input, the command-prompt brackets are empty.

For example:

```
Please create a fully-qualified hostname for this Gateway

(Ex: "mail3.example.com"):
[]> mail3.example.com
```

When there is a default setting, the setting is displayed within the command-prompt brackets. For example:

```
Ethernet interface:
1. Data 1
2. Data 2
3. Management
[1]> 1
```

When a default setting is shown, typing Return is equivalent to typing the default:

```
Ethernet interface:
1. Data 1
2. Data 2
3. Management
[1]> (type Return)
```

Command Syntax

When operating in the interactive mode, the CLI command syntax consists of single commands with no white spaces and no arguments or parameters. For example:

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

Select Lists

When you are presented with multiple choices for input, some commands use numbered lists. Enter the number of the selection at the prompt.

For example:

```
Log level:
1. Error
2. Warning
3. Information
4. Debug
5. Trace
[3]> 3
```

Yes/No Queries

When given a yes or no option, the question is posed with a default in brackets. You may answer **Y**, **N**, **Yes**, or **No**. Case is not significant.

For example:

```
Do you want to enable FTP on this interface?  [Y]> n
```

Subcommands

Some commands give you the opportunity to use subcommands. Subcommands include directives such as NEW, EDIT, and DELETE. For the EDIT and DELETE functions, these commands provide a list of the records previously configured in the system.

For example:

```
mail3.example.com> interfaceconfig



Currently configured interfaces:

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



Choose the operation you want to perform:

- NEW - Create a new interface.

- EDIT - Modify an interface.
```

```
- GROUPS - Define interface groups.

- DELETE - Remove an interface.

[]>
```

Within subcommands, typing Enter or Return at an empty prompt returns you to the main command.

Escape

You can use the Control-C keyboard shortcut at any time within a subcommand to immediately exit return to the top level of the CLI.

History

The CLI keeps a history of all commands you type during a session. Use the Up and Down arrow keys on your keyboard, or the Control-P and Control-N key combinations, to scroll through a running list of the recently-used commands.

```
mail3.example.com> (type the Up arrow key)


mail3.example.com> interfaceconfig (type the Up arrow key)


mail3.example.com> topin (type the Down arrow key)
```

Command Completion

The Cisco IronPort AsyncOS CLI supports command completion. You can type the first few letters of some commands followed by the Tab key, and the CLI completes the string for unique commands. If the letters you entered are not unique among commands, the CLI "narrows" the set. For example:

```
mail3.example.com> set (type the Tab key)
setgateway, sethostname, settime, settz
mail3.example.com> seth (typing the Tab again completes the entry with sethostname)
```

For both the history and file completion features of the CLI, you must type Enter or Return to invoke the command.

Configuration Changes

You can make configuration changes to Cisco IronPort AsyncOS while email operations proceed normally.

Configuration changes will not take effect until you:

1. Issue the `commit` command at the command prompt.
2. Give the `commit` command the input required.

3. Receive confirmation of the `commit` procedure at the CLI.

Changes to configuration that have not been committed will be recorded but not put into effect until the `commit` command is run.

Note Not all commands in AsyncOS require the `commit` command to be run. See Appendix A, "AsyncOS Quick Reference Guide," in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* or view the *Cisco IronPort AsyncOS CLI Reference Guide* for a summary of commands that require commit to be run before their changes take effect.

Exiting the CLI session, system shutdown, reboot, failure, or issuing the `clear` command clears changes that have not yet been committed.

General Purpose CLI Commands

This section describes the commands used to commit or clear changes, to get help, and to quit the command-line interface.

Committing Configuration Changes

The `commit` command is critical to saving configuration changes to the Cisco IronPort appliance. Many configuration changes are not effective until you enter the `commit` command. (A few commands do not require you to use the `commit` command for changes to take effect. See Appendix A, "AsyncOS Quick Reference Guide," in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information. The `commit` command applies configuration changes made to Cisco IronPort AsyncOS since the last `commit` command or the last `clear` command was issued. You may include comments up to 255 characters. Changes are not verified as committed until you receive confirmation along with a timestamp.

Entering comments after the commit command is optional.

```
mail3.example.com> commit


Please enter some comments describing your changes:


[]> Changed "psinet" IP Interface to a different IP address

Changes committed: Wed Jan 01 12:00:01 2003
```

Note To successfully commit changes, you must be at the top-level command prompt. Type Return at an empty prompt to move up one level in the command line hierarchy.

Clearing Configuration Changes

The `clear` command clears any changes made to the Cisco IronPort AsyncOS configuration since the last `commit` or `clear` command was issued.

```
mail3.example.com> clear


Are you sure you want to clear all changes since the last commit? [Y]> y


Changes cleared: Mon Jan 01 12:00:01 2003

mail3.example.com>
```

Quitting the Command Line Interface Session

The `quit` command logs you out of the CLI application. Configuration changes that have not been committed are cleared. The `quit` command has no effect on email operations. Logout is logged into the log files. (Typing `exit` is the same as typing quit.)

```
mail3.example.com> quit


Configuration changes entered but not committed.  Exiting will lose changes.

Type 'commit' at the command prompt to commit changes.

Are you sure you wish to exit?  [N]> Y
```

Seeking Help on the Command Line Interface

The `help` command lists all available CLI commands and gives a brief description of each command. The `help` command can be invoked by typing either `help` or a single question mark (`?`) at the command prompt.

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

CHAPTER 3

Setup and Installation

This chapter guides you through the process of configuring your Cisco IronPort C- or X-Series appliance for email delivery using the System Setup Wizard. If you are configuring an Cisco IronPort M-Series appliance, please see Chapter 17, "The Cisco IronPort M-Series Security Management Appliance". When you have completed this chapter, the Cisco IronPort appliance will be able to send SMTP email over the Internet or within your network.

To configure your system as an Enterprise Gateway (accepting email from the Internet), complete this chapter first, and then see Chapter 5, "Configuring the Gateway to Receive Email" for more information.

Installation Planning

Before You Begin

You can install your Cisco IronPort appliance into your existing network infrastructure in several ways. This section addresses several options available to you as you plan your installation.

Plan to Place the Cisco IronPort Appliance at the Perimeter of Your Network

Please note that your Cisco IronPort appliance is designed to serve as your SMTP gateway, also known as a mail exchanger or "MX." In addition to the "hardened" operating system dedicated for Internet messaging, many of the newest features in the AsyncOS operating system function optimally when the appliance is situated at 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. For example:

- The per-recipient reputation filtering, anti-spam, anti-virus, and Virus Outbreak Filter features (see Reputation Filtering, page 7-1, Cisco IronPort Anti-Spam Filtering, page 9-4, Sophos Anti-Virus Filtering, page 8-2, and Outbreak Filters, page 10-1) are designed to work with a *direct flow* of messages from the Internet and from your internal network. You can configure the Cisco IronPort appliance for policy enforcement (The Host Access Table (HAT): Sender Groups and Mail Flow Policies, page 5-7) for all email traffic to and from your enterprise.

You need to ensure that the Cisco IronPort appliance 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 Cisco IronPort appliance 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 SenderBase Reputation Service for the sender's SenderBase Reputation Score (SBRS), and to improve the efficacy of the Cisco IronPort Anti-Spam and Outbreak Filters features.

Note If you cannot configure the appliance as the *first* machine receiving email from the Internet, you can still exercise some of the security services available on the appliance. Refer to Incoming Relays, page 9-19 for more information.

When you use the Cisco IronPort appliance as your SMTP gateway:

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

Register the Cisco IronPort Appliance in DNS

Malicious email senders actively search public DNS records to hunt for new victims. You need to ensure that the Cisco IronPort appliance is registered in DNS, if you want to utilize the full capabilities of Cisco IronPort Anti-Spam, Outbreak Filters, McAfee Antivirus and Sophos Anti-Virus. To register the Cisco IronPort appliance in DNS, create an A record that maps the appliance's hostname to its IP address, and an MX record that maps your public domain to the appliance's hostname. You must specify a priority for the MX record to advertise the Cisco IronPort appliance as either a primary or backup MTA for your domain.

In the following example, the Cisco IronPort appliance (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 Cisco IronPort appliance 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 Cisco IronPort appliance to have an MX record priority of equal or higher value than the rest of your MTAs.

Installation Scenarios

You may want to review all features of the appliance prior to installing. Chapter 4, "Understanding the Email Pipeline" provides an overview of all functions available on the appliance that may affect the placement of the Cisco IronPort appliance within your infrastructure.

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 IronPort Customer Support (see Cisco IronPort Support Community, page 1-10).

Configuration Overview

n an enterprise

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

The following network scenarios are described:

- Behind the Firewall (see Figure 3-2 on page 3-7)

Choose the configuration that best matches your infrastructure. Then proceed to the next section, Preparing for Setup, page 3-8.

Incoming

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

Outgoing

- Outgoing mail sent by internal users is routed by the groupware server to the Cisco IronPort appliance.
- The Cisco IronPort appliance accepts outbound email based on settings in the Host Access Table for the private listener. (For more information, see Receiving Email with Listeners, page 5-1.)

Ethernet Interfaces

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

See "Using Virtual Gateway™ Technology" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* and Appendix B, "Assigning Network and IP Addresses" for more information about assigning multiple IP addresses to the available interfaces.

Note The Cisco IronPort X1000/1050/1060/1070, C60/600/650/660/670, and C30/300/350/360/370 Email Security appliances have three available Ethernet interfaces by default. The Cisco IronPort C10/100/150/160 Email Security appliances have two available Ethernet interfaces.

Advanced Configurations

In addition to this configurations shown in Figure 3-2 and Figure 3-3, you can also configure:

- Multiple Cisco IronPort appliances using the Centralized Management feature
- Redundancy at the network interface card level by "teaming" two of the Ethernet interfaces on Cisco IronPort appliances using the NIC Pairing feature.

Both of these features are discussed in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Firewall Settings (NAT, Ports)

Depending on your network configuration, your firewall may need to be configured to allow access on the following ports.

SMTP and DNS services must have access to the Internet. For other system functions, the following services may be required:

Table 3-1 Firewall Ports

• SMTP: port 25	• LDAP: port 389 or 3268
• DNS: port 53	• NTP: port 123
• HTTP: port 80	• LDAP over SSL: port 636
• HTTPS: port 443	• LDAP with SSL for Global Catalog queries: port 3269
• SSH: port 22	• FTP: port 21, data port TCP 1024 and higher
• Telnet: port 23	• Cisco IronPort Spam Quarantine: port 6025

Appendix C, "Firewall Information" contains all information about the possible ports that may need to be opened for proper operation of the Cisco IronPort appliance. For example, ports in the firewall may need to be opened for connections:

- from the external clients (MTAs) to the Cisco IronPort appliance
- to and from groupware servers
- to the Internet root DNS servers or internal DNS servers
- to the Cisco IronPort downloads servers for McAfee and Sophos Anti-Virus updates, Outbreak Filters rules, and updates to AsyncOS
- to the NTP servers
- to LDAP servers

Support Languages

AsyncOS can display its GUI and CLI in any of the following languages:

- English
- French
- Spanish
- German
- Italian
- Korean
- Japanese
- Portuguese (Brazil)
- Chinese (zh-cn and zh-tw)
- Russian

Physical Dimensions

The following physical dimensions apply to the **Cisco IronPort X1050/1060, C650/660,** and **C350/360** Email Security appliances:

- Height: 8.656 cm (3.40 inches)
- Width: 48.26 cm (19.0 inches) with rails installed (without rails, 17.5 inches)
- Depth: 75.68cm (29.79 inches)
- Weight: maximum 26.76 kg (59 lbs)

The following physical dimensions apply to the **Cisco IronPort X1070, C670 and C370** Email Security appliances:

- Height: 8.64 cm (3.40 inches)
- Width: 48.24 cm (18.99 inches) with or without rails
- Depth: 72.06 cm (28.40 inches)
- Weight: maximum 23.59 kg (52 lbs)

The following physical dimensions apply to the **Cisco IronPort C150 and C160** Email Security appliances:

- Height: 4.2 cm (1.68 inches)
- Width: 48.26 cm (19.0 inches) with rails installed (without rails, 17.5 inches)
- Depth: 57.6 cm (22.7 inches)
- Weight: maximum 11.8 kg (26 lbs)

Physically Connecting the Cisco IronPort Appliance to the Network

Configuration Scenarios

The typical configuration scenario for the Cisco IronPort appliance is as follows:

- **Interfaces** - Only one of the three available Ethernet interfaces on the Cisco IronPort appliance 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 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 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 IronPort C- or X-Series appliance.
  - The Cisco IronPort appliance 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.

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 appliance 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, page 3-10). Most configuration scenarios are represented by one of the following three figures.

Figure 3-2 Behind the Firewall Scenario / 2 Listeners Configuration

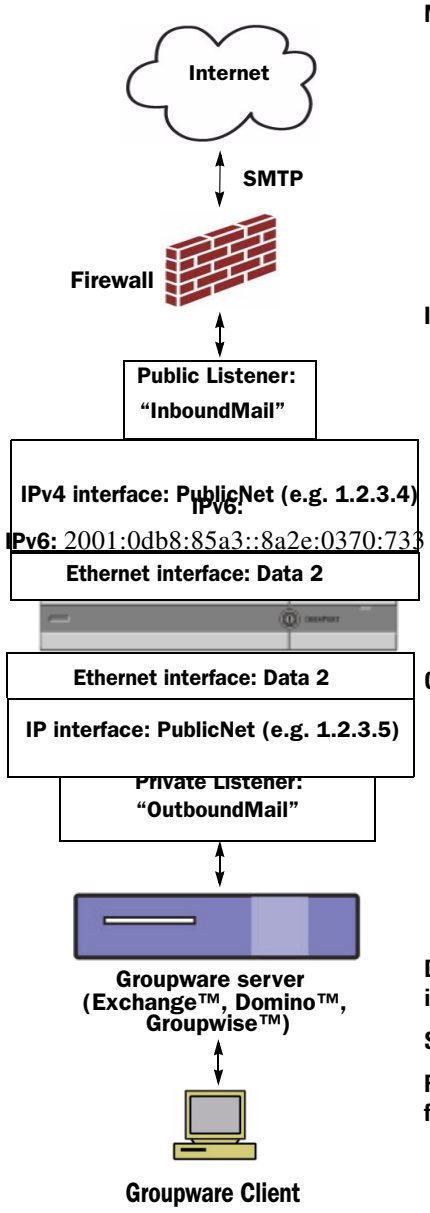

Notes:

- 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 Cisco IronPort appliance

Figure 3-3 One Listener Configuration

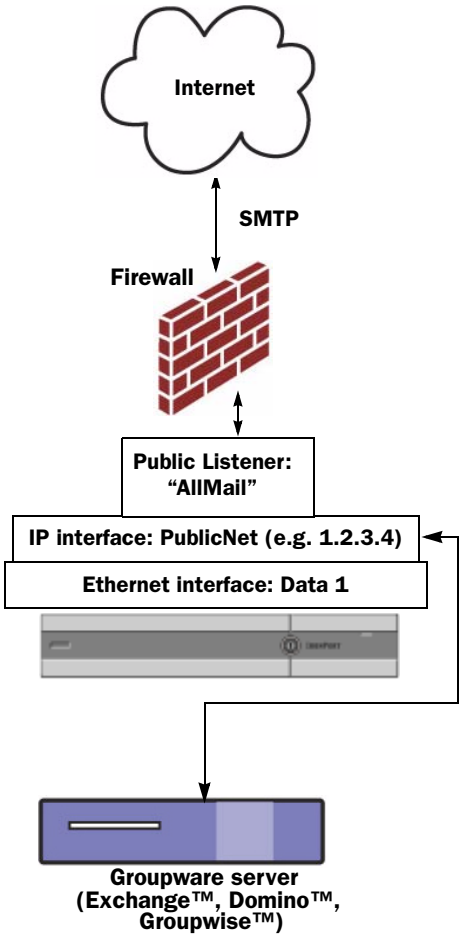

Notes:

- 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 Cisco IronPort appliance

Preparing for Setup

Step 1 Determine how you will connect to the appliance.

Step 2 Determine network and IP address assignments. You can use both IPv4 and IPv6 addresses.

Step 3 Gather information about your system setup.

Step 4 Launch a web browser and enter the IP address of the appliance. (Alternatively, you can access the command line interface (CLI) described in Running the Command Line Interface (CLI) System Setup Wizard, page 3-26)

Step 5 Run the System Setup Wizard to configure your system.

Determine Method for Connecting to the Appliance

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

Connecting to the Appliance

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

Table 3-2 Options for Connecting to the Appliance

Ethernet	An Ethernet connection between a PC and the network and between the network and the Cisco IronPort 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 Cisco IronPort Serial Console port. If you cannot use the Ethernet method, a straight serial-to-serial connection between the computer and the appliance will work until alternate network settings can be applied to the Management port. For pinout information, see Accessing via a Serial Connection, page A-7. 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 Appendix A, "Accessing the Appliance" for more information.) You can also create multiple user accounts with differing administrative privileges to access the appliance. (For more information, see "Adding Users" in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.)

Determining Network and IP Address Assignments

Choosing Network Connections to Receive and Deliver Email

Most users take advantage of the two Data Ethernet ports on the Cisco IronPort appliance by connecting to two networks from the appliance:

- 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 appliance 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 Security appliance 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 IronPort's time servers)

See Appendix B, "Assigning Network and IP Addresses" for more information.

Note If you are running a firewall on your network between the Internet and the Cisco IronPort appliance, it may be necessary to open specific ports for the Cisco IronPort appliance to work properly. See Appendix C, "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 Appendix B, "Assigning Network and IP Addresses" for more detailed information on network and IP addresses. See Chapter 17, "The Cisco IronPort M-Series Security Management Appliance" if you are configuring a Cisco IronPort M-Series appliance.

Table 3-3 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 Password:		
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:		
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**	

Table 3-3 System Setup Worksheet: 2 Listeners for Segregating Email Traffic (Continued)

Message Security	
SenderBase 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 3-4 System Setup Worksheet: 1 Listener for All Email Traffic

System Settings		
Default System Hostname:		
Email System Alerts To:		
Deliver Scheduled Reports To:		
Time Zone:		
NTP Server:		
Admin Password:		
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		
SenderBase Reputation Filtering:	Enable / Disable	
Anti-Spam Scanning Engine	None / IronPort	
McAfee Anti-Virus Scanning Engine	Enable / Disable	

Table 3-4 System Setup Worksheet: 1 Listener for All Email Traffic (Continued)

Sophos Anti-Virus Scanning Engine	Enable / Disable
Outbreak Filters	Enable / Disable

Using the System Setup Wizard

The Cisco IronPort AsyncOS operating system provides a browser-based System Setup Wizard to guide you through the five step process of system configuration. Also included is a command line interface (CLI) version of the System Setup Wizard. See Running the Command Line Interface (CLI) System Setup Wizard, page 3-26 for more information. Some users will want to take advantage of custom configuration options not available in the System Setup Wizard. However, you must use the System Setup Wizard for the initial setup to ensure a complete configuration. If you have gathered the information required in Preparing for Setup, page 3-8, the configuration process will take less time to complete.

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

Warning **The Cisco IronPort appliance ships with a default IP address of `192.168.42.42` on the Management port of C650/660/670, C350/360/370, and X1050/1060/1070 systems, and the Data 1 port of C150/160 systems. Before connecting the Cisco IronPort appliance to your network, ensure that no other device's IP address conflicts with this factory default setting. If you are configuring a Cisco IronPort M-Series appliance, please see The Cisco IronPort M-Series Security Management Appliance, page 17-1.**

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

Accessing the Web-Based Graphical User Interface (GUI)

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

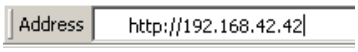

The login screen is displayed:

Figure 3-4 Logging in to the Appliance

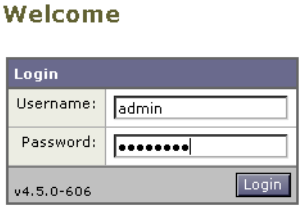

Log in to the appliance by entering the username and password below.

Factory Default Username and Password

- Username: `admin`
- Password: `ironport`

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

Running the Web-Based System Setup Wizard

To launch the System Setup Wizard, log in to the graphical user interface as described in Accessing the Web-Based Graphical User Interface (GUI), page 3-13. On the System Administration tab, click System Setup Wizard in the list of links on the left. On brand new (not upgraded from previous releases of AsyncOS) systems, your browser will automatically be redirected to the System Setup Wizard.

The System Setup Wizard walks you through completing the following configuration tasks, broken down into 5 categories:

Step 1 Start

- Read and accept the license agreement

Step 2 System

- Setting the hostname of the appliance
- Configuring alert settings, report delivery settings, and AutoSupport
- Setting the system time settings, and NTP server
- Resetting the admin password
- Enabling SenderBase Network participation

Step 3 Network

- 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 appliance (optional)

Step 4 Security

- Enabling SenderBase Reputation Filtering
- Enabling the Anti-Spam service
- Enabling the Cisco IronPort Spam Quarantine
- Enabling the Anti-Virus service
- Enabling the Outbreak Filters service

Step 5 Review

- Reviewing your setup and installing the configuration

Step through the System Setup Wizard, clicking **Next** after you complete each step. You can move back to a previous step by clicking **Previous**. At the end of the process, you are prompted to commit the changes you have made. Your changes will not take effect until they have been committed. If you click **Next**, but have left a required field blank (or entered incorrect information), the fields in question are outlined in red. Make your corrections and click **Next** again.

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.

Figure 3-5 System Setup Wizard: Step 1. Start

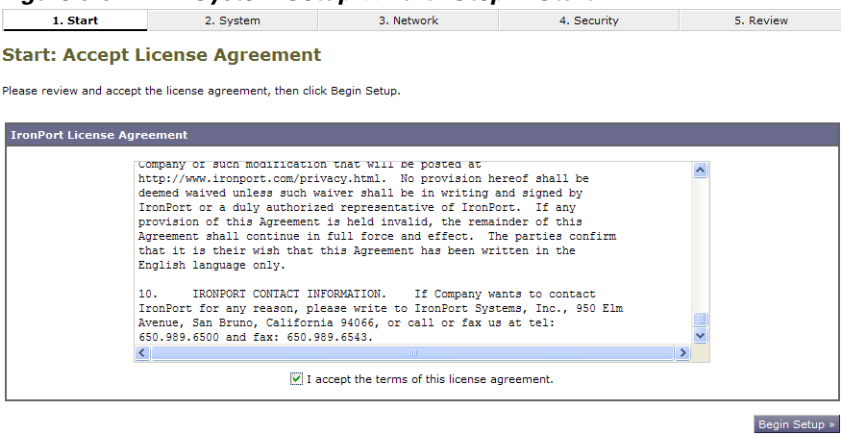

You can also view the text of the agreement here:

`https://support.ironport.com/license/eula.html`

Step 2: System

Setting the Hostname

Define the fully-qualified hostname for the Cisco IronPort appliance. This name should be assigned by your network administrator.

Configuring System Alerts

Cisco IronPort 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, page 15-15.

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 appliance rather than delivered.

Setting the Time

Set the time zone on the Cisco IronPort appliance 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, page 15-48 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 IronPort Systems time servers (`time.ironport.com`) to synchronize the time on your Cisco IronPort appliance is already configured.

Setting the Password

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

Participating in SenderBase Network

SenderBase is an email reputation service designed to help email administrators research senders, identify legitimate sources of email, and block spammers.

If you agree to participate in the SenderBase Network, Cisco will collect aggregated email traffic statistics about your organization. This includes only summary data on message attributes and information on how different types of messages were handled by Cisco IronPort appliances. For example, Cisco does not collect the message body or the message subject. Personally identifiable information or information that identifies your organization will be kept confidential. To learn more about SenderBase, including examples of the data collected, follow the **Click here for more information about what data is being shared…** link (see Frequently Asked Questions, page 13-2).

To participate in the SenderBase Network, check the box next to “Allow IronPort to gather anonymous statistics on email and report them to SenderBase in order to identify and stop email-based threats” and click **Accept**.

See Chapter 13, “SenderBase Network Participation” for more information.

Enabling AutoSupport

The Cisco IronPort AutoSupport feature (enabled by default) keeps the Cisco IronPort Customer Support team aware of issues with your Cisco IronPort appliance so that we can provide better support to you. (For more information, see Cisco IronPort AutoSupport, page 15-17.)

Figure 3-6 System Setup Wizard: Step 2. System

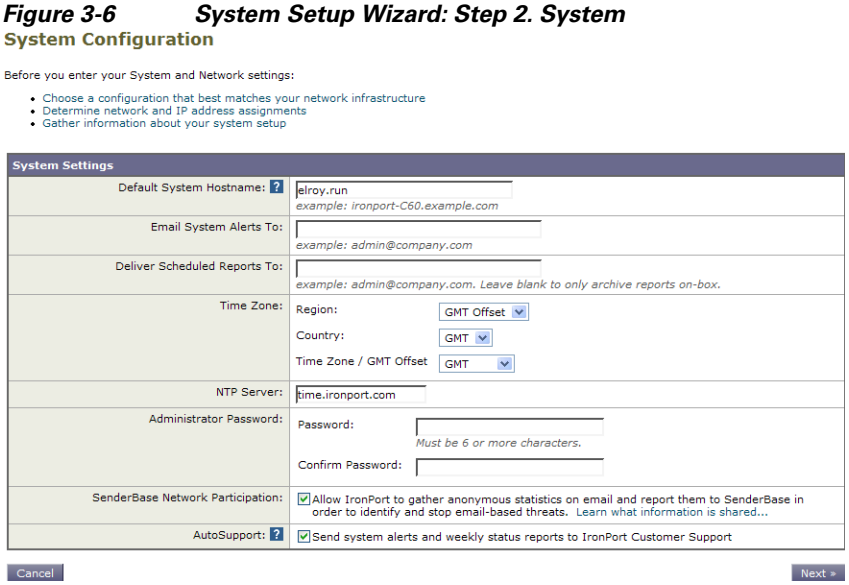

Click **Next** to continue.

Step 3: Network

In Step 3, you define the default router (gateway) and configure the DNS settings, and then set up the appliance 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 IronPort 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, page 15-39.

Note The appliance 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 appliance while you are setting up the appliance, 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 Cisco IronPort appliance has network interfaces that are associated with the physical ports on the machine. For example, on C650/660/670, C350/360/370, and X1050/1060/1070 appliances, three physical Ethernet interfaces are available. On C150/160 appliances, two physical Ethernet interfaces are available.

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. Typically, you would use one interface for incoming, one for outgoing, and one for appliance management. On the C150 and C160 appliances, 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 appliance. If you decide to use both the Data 1 Ethernet port and the Data 2 Ethernet port, you need this information for both connections.

C650/660/670, C350/360/370, and X1050/1060/1070 customers: 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.

C150/160 customers: 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, page 3-10.

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 Appendix B, "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" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information.)

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 Recipients, page 5-51.

Relaying Mail (Optional)

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

These are entries in the RELAYLIST of the Host Access Table for a listener. See Sender Group Syntax, page 5-21 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 appliance.

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.

Note The following example pertains to X1050/1060/1070, C650/660/670, and C350/360/370 appliances. For C150/160 appliances, the Data 2 interface is typically configured for both incoming and outgoing mail while the Data 1 interface is used for appliance management (see C150/160 Installations, page 3-20).

Figure 3-7 Network Interfaces: 2 Interfaces in Addition to Management (Segregated Traffic)

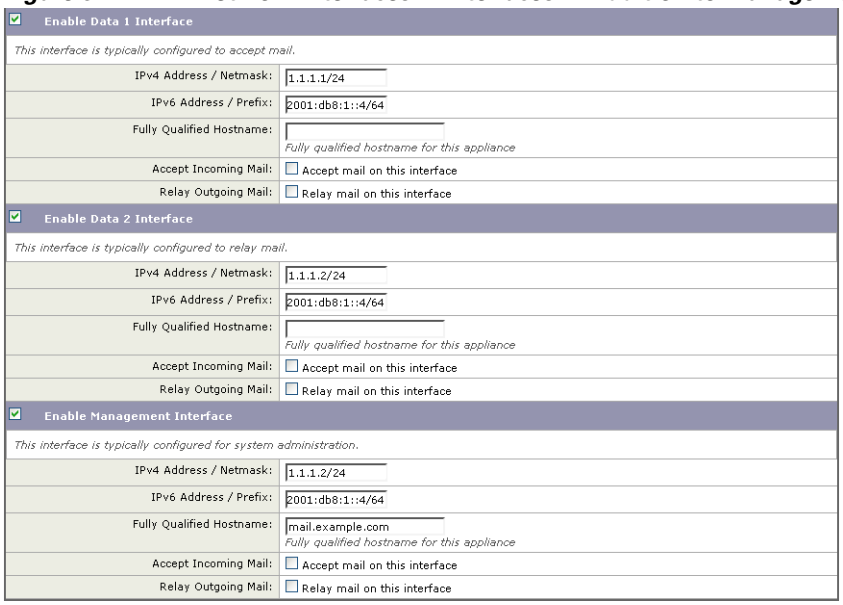

Use this configuration if you want your network to look like Figure 3-2 on page 3-7.

C150/160 Installations

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

Figure 3-8 Network Interfaces: 1 IP Address for Incoming and Outgoing (Nonsegregated) Traffic

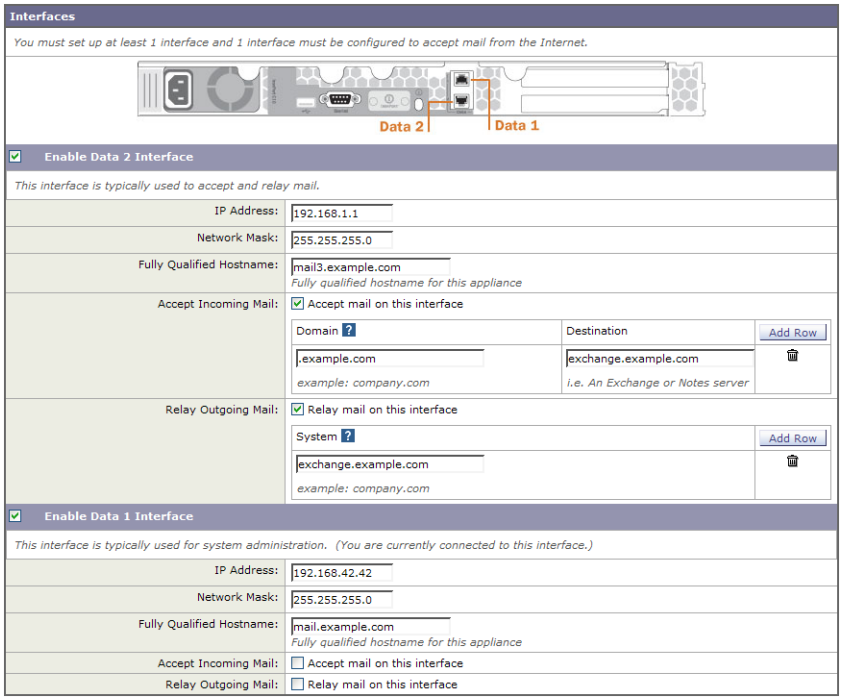

Use this configuration if you want your network to look like Figure 3-3 on page 3-8.

Click **Next** to continue.

Step 4: Security

In step 4, you configure anti-spam and anti-virus settings. The anti-spam options include SenderBase 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 SenderBase Reputation Filtering

The SenderBase 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 Cisco IronPort Anti-Spam.

The SenderBase Reputation Service (http://www.senderbase.org) provides an accurate, flexible way for users to reject or throttle suspected spam based on the connecting IP address of the remote host. The SenderBase Reputation Service returns a score based on the probability that a message from a given source is spam. The SenderBase 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 SenderBase Reputation Filtering.

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

Enabling Anti-Spam Scanning

Your Cisco IronPort appliance may ship with a 30-day evaluation key for Cisco IronPort Anti-Spam software. During this portion of the System Setup Wizard, you can choose to enable Cisco IronPort Anti-Spam globally on the appliance. 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 Cisco IronPort Spam Quarantine. The Cisco IronPort Spam Quarantine serves as the end-user quarantine for the appliance. Only administrators can access the quarantine until end-user access is configured.

See Chapter 9, “Anti-Spam” for all of the Cisco IronPort Anti-Spam configuration options available on the appliance. See “Quarantines” in the *Cisco IronPort AsyncOS for Email Daily Management Guide* for information about the Cisco IronPort Spam Quarantine.

Enabling Anti-Virus Scanning

Your Cisco IronPort appliance 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 appliance.

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

See Chapter 8, “Anti-Virus” for all of the anti-virus configuration options available on the appliance.

Enabling Outbreak Filters

Your Cisco IronPort appliance 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 Chapter 10, “Outbreak Filters” for more information.

Figure 3-9 System Setup Wizard: Step 4. Configuring Message Security

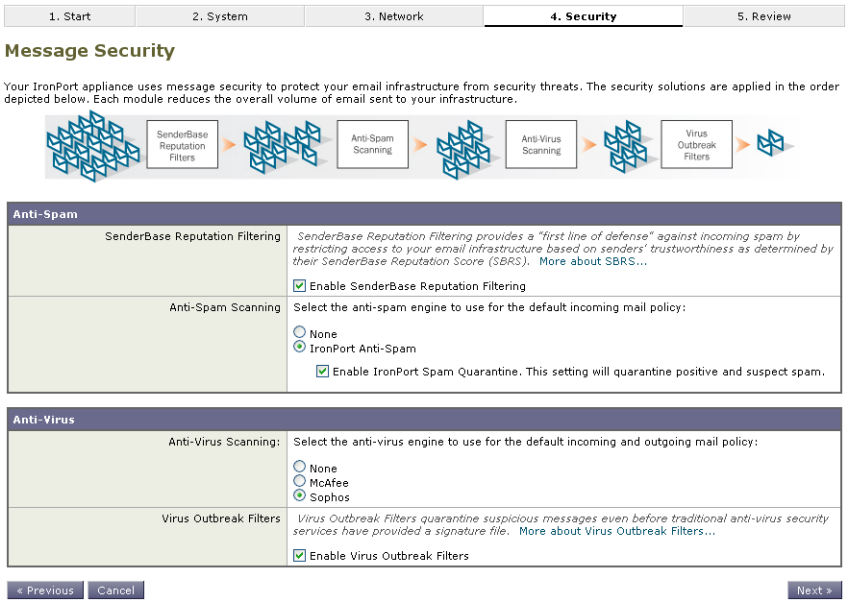

Click **Next** to continue.

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.

Figure 3-10 System Setup Wizard: Step 5. Review

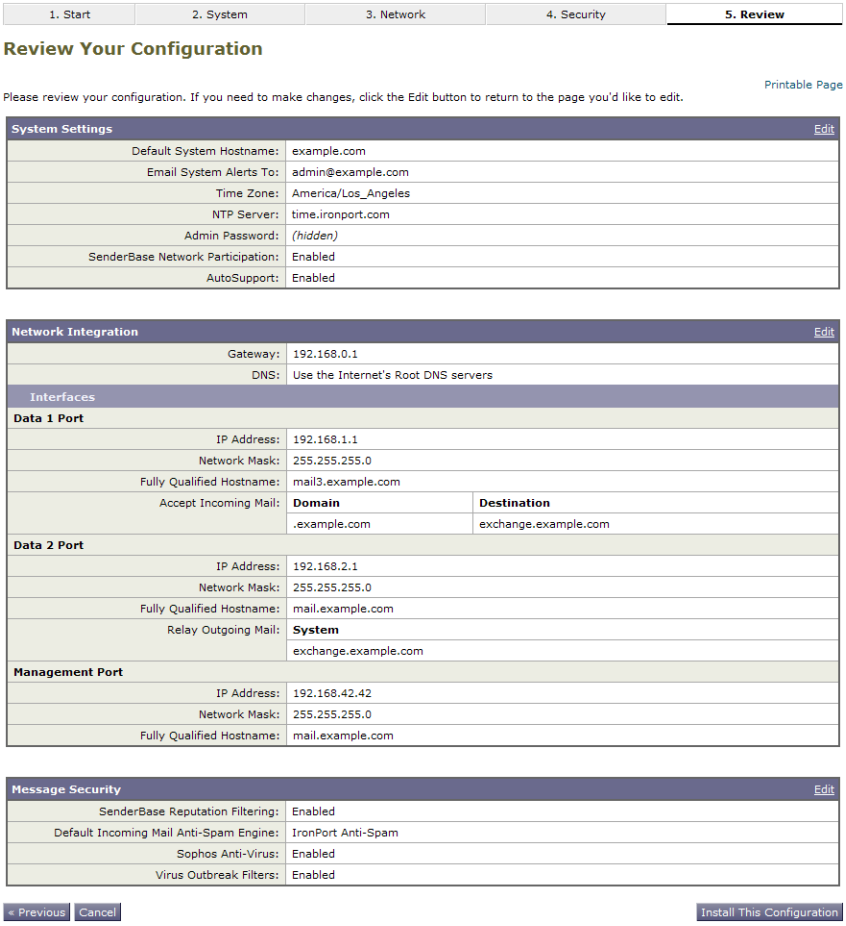

Once you are satisfied with the information displayed click **Install This Configuration**. A confirmation dialog is displayed. Click **Install** to install the new configuration.

Figure 3-11 System Setup Wizard: Confirm Install

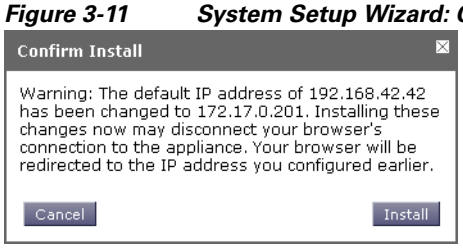

Your Cisco IronPort appliance is now ready to send email.

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 appliance (the Management interface on X1050/1060/1070, C650/660/670, and C350/360/370 systems, or the Data 1 interface on C150/160 systems) 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, page 3-38 for more information.

Configuring Active Directory

If the System Setup Wizard properly installs the configuration on the Email Security appliance, 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.

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

Figure 3-12 Active Directory Wizard – Step 1: Start

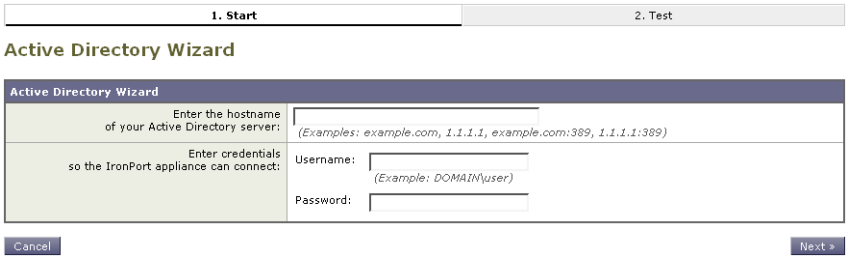

Step 2 Enter the host name for the Active Directory server.

Step 3 Enter a username and password 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.

Figure 3-13 Active Directory Wizard – Step 2: Directory Lookup Test

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 appliance to work with your Active Directory Wizard, or skip the process, the System Setup Next Steps page appears.

Figure 3-14 System Setup is Complete

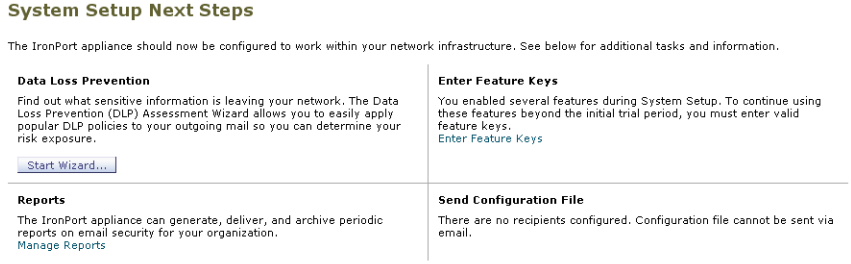

Click the links on the System Setup Next Steps page to proceed with the configuration of your Cisco IronPort appliances.

Accessing the Command Line Interface (CLI)

Access to the CLI varies depending on the management connection method you chose in Connecting to the Appliance, page 3-9. The factory default username and password 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 "Common Administrative Tasks" in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.) The System Setup Wizard asks you to change the password for the admin account. The password for the admin account can also be reset directly at any time using the password command.

To connect via Ethernet: Start an SSH or Telnet session with the factory default IP address 192.168.42.42. SSH is configured to use port 22. Telnet is configured to use port 23. Enter the username and password 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 Appliance, page 3-9. Enter the username and password below.

Log in to the appliance by entering the username and password below.

Factory Default Username and Password

- Username: **admin**
- Password: **ironport**

For example:

```
login: admin
password: ironport
```

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 appliance, 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 Running the Web-Based System Setup Wizard, page 3-14.

Change the Admin Password

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

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 Cisco IronPort appliance. 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 X1000/1050/1060/1070, C60/600/650/660/670, and C30/300/350/360/370 appliances) or Data 1 (on C10/100/150/160 appliances), and then prompts you to configure a logical IP interface on any other physical Ethernet interfaces available on the appliance.

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.

X1050/1060/1070, C650/660/670, and C350/360/370 customers: 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.

C150/160 customers: 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 address** 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 Appendix B, “Assigning Network and IP Addresses”for more detailed information on Network and IP Address configuration.

Note For C10/100 customers, 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 appliance (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. For more information, see "Enabling a Certificate for HTTPS" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Configure the DNS Settings

Next, you configure the DNS (Domain Name Service) settings. Cisco IronPort 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 Cisco IronPort appliance — either from your internal systems or from the Internet. Cisco IronPort 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.

X1050/1060/1070, C650/660/670 and C350/360/370 customers: 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, page 5-1.)

C150/160 customers: 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 C10/100/150/160 Listener Example, page 3-32.

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" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information.)
- Whether or not to enable filtering based on SenderBase Reputation Scores (SBRS) 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 appliance (private listeners). (These are the first Recipient Access Table and Host Access Table entries for a listener. See Sender Group Syntax, page 5-21 and Accepting Email for Local Domains or Specific Users on Public Listeners (RAT), page 5-50 for more information.)

Public Listener

Note The following examples of creating a public and private listener apply to X1050/1060/1070, C650/660/670, and C350/360/370 customers only. Cisco IronPort C150/160 customers should skip to the next section C10/100/150/160 Listener Example, page 3-32.

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 Cisco IronPort appliance 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, page 5-21.

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

```
You are now going to configure how the IronPort C60 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 Public and Private Listeners, page 5-3.

```
Do you want to configure the C60 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 IronPort C60.

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.

*****
```

C10/100/150/160 Listener Example

Note The following example of creating a listener applies to C150/160 customers 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 Cisco IronPort appliance 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, page 5-21.

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

```
You are now going to configure how the IronPort C10 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 IronPort C10.

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 C10/100 customers, all outgoing mail will be calculated in the Mail Flow Monitor feature (which is normally used for inbound messages). See "Using the Email Security Monitor" in the *Cisco IronPort AsyncOS for Email Daily Management Guide* for more information.

Enable Cisco IronPort Anti-Spam

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

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

Note If you do not accept the license agreement, Cisco IronPort Anti-Spam is not enabled on the appliance.

See Chapter 9, "Anti-Spam" for all of the Cisco IronPort Anti-Spam configuration options available on the appliance.

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 Cisco IronPort 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 Cisco IronPort Spam Quarantine. Enabling the Cisco IronPort Spam Quarantine also enables the end-user quarantine on the appliance. Only administrators can access the end-user quarantine until end-user access is configured.

See "Quarantines" in the *Cisco IronPort AsyncOS for Email Daily Management Guide* for information on the Cisco IronPort Spam Quarantine.

Enable Anti-Virus Scanning

Your Cisco IronPort appliance 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 appliance. You must accept a license agreement for each anti-virus scanning engine you want to enable on your appliance.

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

See Chapter 8, "Anti-Virus" for the anti-virus configuration options available on the appliance.

Enable Outbreak FiltersOutbreak Filters and SenderBase Email Traffic Monitoring Network

This next step prompts you to enable both SenderBase participation and Outbreak Filters. Your Cisco IronPort appliance ships with a 30-day evaluation key for Outbreak Filters.

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, page 10-1.

SenderBase Participation

SenderBase is an email reputation service designed to help email administrators research senders, identify legitimate sources of email, and block spammers.

If you agree to participate in the SenderBase Email Traffic Monitoring Network, Cisco will collect aggregated statistics about email sent to your organization. This includes summary data on message attributes and information on how different types of messages were handled by Cisco IronPort appliances.

See Chapter 13, "SenderBase Network Participation" for more information.

Configure the Alert Settings and AutoSupport

Cisco IronPort 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 Alerts, page 15-15.

The Cisco IronPort AutoSupport feature keeps the Cisco IronPort Customer Support team aware of issues with your Cisco IronPort appliance 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 Cisco IronPort AutoSupport, page 15-17.)

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 appliance instead of sent via email.

Configure Time Settings

Cisco IronPort 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 Cisco IronPort appliance so that timestamps in message headers and log files are correct. You can also use the Cisco IronPort Systems time servers to synchronize the time on your Cisco IronPort appliance.

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 Cisco IronPort appliance is now ready to send email.

Test the Configuration

To test the Cisco IronPort 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 Cisco IronPort appliance uses feature keys to enable features. The first time you create a listener in the `systemsetup` command, enable Cisco IronPort 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, page 3-15.

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 IronPort 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 the section on working with feature keys in "Common Administrative Tasks" in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.)

What's Next: Understanding the Email Pipeline

Now that systemsetup is complete, your Cisco IronPort appliance 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 appliances' configuration. Chapter 4, "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.

CHAPTER 4

Understanding the Email Pipeline

The Email Pipeline is the process or flow email follows as it is processed by the Cisco IronPort appliance. Understanding the Email Pipeline is essential to getting the best performance from your Cisco IronPort appliance.

This chapter provides an overview of the Email Pipeline for incoming mail and a brief description of each feature. The brief description also includes a link to the chapter or book containing a detailed explanation of the feature.

Overview: Email Pipeline

Table 4-1 and Table 4-2 provide an overview of how email is processed through the system, from reception to routing to delivery. Each feature is processed in order (from top to bottom) and is briefly described below. A full description of each feature can be found in the following chapters. Some features are described in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* and *Cisco IronPort AsyncOS for Email Daily Management Guide*.

Shaded areas in Table 4-2 represent processing that occurs in the Work Queue (see Work Queue / Routing, page 4-6). You can test most of the configurations of features in this pipeline using the `trace` command. For more information, seeDebugging Mail Flow Using Test Messages: Trace, page -446.

Table 4-1 Email Pipeline for the Cisco IronPort Appliance: Receiving Email Features

Feature	Description
Host Access Table (HAT) **Host DNS Sender Verification** **Sender Groups** **Envelope Sender Verification** **Sender Verification Exception Table** **Mail Flow Policies**	ACCEPT, REJECT, RELAY, or TCPREFUSE connections Maximum outbound connections Maximum concurrent inbound connections per IP address Maximum message size and messages per connection Maximum recipients per message and per hour TCP listen queue size TLS: no/preferred/required SMTP AUTH: no/preferred/required Drop email with malformed FROM headers Always accept or reject mail from entries in the Sender Verification Exception Table. SenderBase on/off (IP profiling/flow control)
Received Header	Adds a received header to accepted email: on/off.
Default Domain	Adds default domain for “bare” user addresses.
Bounce Verification	Used to verify incoming bounce messages as legitimate.
Domain Map	Rewrites the Envelope Recipient for each recipient in a message that matches a domain in the domain map table.
Recipient Access Table (RAT)	(Public listeners only) ACCEPT or REJECT recipients in `RCPT TO` plus Custom SMTP Response. Allow special recipients to bypass throttling.
Alias tables	Rewrites the Envelope Recipient. (Configured system-wide. `aliasconfig` is not a subcommand of `listenerconfig`.)
LDAP Recipient Acceptance	LDAP validation for recipient acceptance occurs within the SMTP conversation. If the recipient is not found in the LDAP directory, the message is dropped or bounced. LDAP validation can be configured to occur within the work queue instead.
SMTP Call-Ahead Validation	SMTP call-ahead recipient validation occurs within the SMTP conversation. The SMTP conversation is paused while the Email Security appliance calls ahead to the external SMTP server. The message is dropped or bounced, or the mailing action is allowed depending on the SMTP server response.

Table 4-2 Email Pipeline for the Cisco IronPort Appliance: Routing and Delivery Features

Work Queue	**LDAP Recipient Acceptance**		LDAP validation for recipient acceptance occurs within the work queue. If the recipient is not found in the LDAP directory, the message is dropped or bounced. LDAP validation can be configured to occur within the SMTP conversation instead.
	Masquerading **or LDAP Masquerading**		Masquerading occurs in the work queue; it rewrites the Envelope Sender, To:, From:, and/or CC: headers, from a static table or via an LDAP query.
	LDAP Routing		LDAP queries are performed for message routing or address rewriting. Group LDAP queries work in conjunction with message filter rules `mail-from-group` and `rcpt-to-group`.
	Message Filters*		Message Filters are applied prior to message "splintering." * Can send messages to quarantines.
	Safelist/Blocklist Scanning		AsyncOS checks the sender address against the end user safelist/blocklist database. If the sender address is safelisted, anti-spam scanning is skipped. The message may be splintered if there are multiple recipients. *Can send messages to quarantines if sender is blocklisted.
	Anti-Spam**	Email Security Manager Scanning (Per Recipient)	Anti-Spam scanning engine examines messages and returns a verdict for further processing.
	Anti-Virus*		Anti-Virus scanning examines messages for viruses. Messages are scanned and optionally repaired, if possible. * Can send messages to quarantines.
	Content Filters*		Content Filters are applied. DKIM, SPF, and SIDF verification is performed if appropriate content filter conditions are defined. * Can send messages to quarantines.
	Outbreak Filters*		The Outbreak Filters feature helps protect against virus outbreaks, as well as new scam, phishing and malware attacks. * Can send messages to quarantines.
	Data Loss Prevention (Outgoing Messages Only)		RSA Email Data Loss prevention examines outgoing messages for sensitive data. RSA Email DLP is for outgoing messages only. * Can send messages to quarantines.
	Virtual gateways		Sends mail over particular IP interfaces or groups of IP interfaces.

Table 4-2 Email Pipeline for the Cisco IronPort Appliance: Routing and Delivery Features

Delivery limits	1. Sets the default delivery interface. 2. Sets the total maximum number of outbound connections.
Domain-based Limits	Defines, per-domain: maximum outbound connections for each virtual gateway and for the entire system; the bounce profile to use; the TLS preference for delivery: no/preferred/required
Domain-based routing	Routes mail based on domain without rewriting Envelope Recipient.
Global unsubscribe	Drops recipients according to specific list (configured system-wide).
Bounce profiles	Undeliverable message handling. Configurable per listener, per Destination Controls entry, and via message filters.

* These features can send messages to special queues called Quarantines.

** Can send messages to the Cisco IronPort Spam Quarantine.

Incoming / Receiving

The receiving phase of the Email Pipeline involves the initial connection from the sender's host. Each message's domains can be set, the recipient is checked, and the message is handed off to the work queue.

Host Access Table (HAT), Sender Groups, and Mail Flow Policies

The HAT allows you to specify hosts that are allowed to connect to a listener (that is, which hosts you will allow to send email).

Sender Groups are used to associate one or more senders into groups, upon which you can apply message filters, and other Mail Flow Policies. Mail Flow Policies are a way of expressing a group of HAT parameters (access rule, followed by rate limit parameters and custom SMTP codes and responses).

Together, sender groups and mail flow policies are defined in a listener's HAT.

Host DNS verification settings for sender groups allow you to classify unverified senders prior to the SMTP conversation and include different types of unverified senders in your various sender groups.

While the connecting host was subject to Host DNS verification in sender groups — prior to the SMTP conversation — the domain portion of the envelope sender is DNS verified in mail flow policies, and the verification takes place during the SMTP conversation. Messages with malformed envelope senders can be ignored. You can add entries to the Sender Verification Exception Table — a list of domains and email addresses from which to accept or reject mail despite envelope sender DNS verification settings.

Reputation Filtering allows you to classify email senders and restrict access to your email infrastructure based on sender's trustworthiness as determined by the Cisco IronPort SenderBase Reputation Service.

For more information, see The Host Access Table (HAT): Sender Groups and Mail Flow Policies, page 5-7.

Received: Header

Using the `listenerconfig` command, you can configure a listener to not include the Received: header by default to all messages received by the listener.

For more information, see "Advanced Configuration Options" in the "Customizing Listeners" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Default Domain

You can configure a listener to automatically append a default domain to sender addresses that do not contain fully-qualified domain names; these are also known as "bare" addresses (such as "joe" vs. "joe@example.com").

For more information, see "SMTP Address Parsing Options" in the "Customizing Listeners" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Bounce Verification

Outgoing mail is tagged with a special key, and so if that mail is sent back as a bounce, the tag is recognized and the mail is delivered. For more information, see "IronPort Bounce Verification" in the "Configuring Routing and Delivery Features" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Domain Map

For each listener you configure, you can construct a domain map table which rewrites the envelope recipient for each recipient in a message that matches a domain in the domain map table. For example, joe@old.com -> joe@new.com

For more information, see "The Domain Map Feature" in the "Configuring Routing and Delivery Features" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Recipient Access Table (RAT)

For inbound email only, the RAT allows you to specify a list of all local domains for which the Cisco IronPort appliance will accept mail.

For more information, see Accepting Email for Local Domains or Specific Users on Public Listeners (RAT), page 5-50.

Alias Tables

Alias tables provide a mechanism to redirect messages to one or more recipients. Aliases are stored in a mapping table. When the envelope recipient (also known as the Envelope To, or `RCPT TO`) of an email matches an alias as defined in an alias table, the envelope recipient address of the email will be rewritten.

For more information about Alias Tables, see "Creating Alias Tables" in the "Configuring Routing and Delivery Features" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

LDAP Recipient Acceptance

You can use your existing LDAP infrastructure to define how the recipient email address of incoming messages (on a public listener) should be handled during the SMTP conversation or within the workqueue. See "Accept Queries" in the "Customizing Listeners" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*. This allows the Cisco IronPort appliance to combat directory harvest attacks (DHAP) in a unique way: the system accepts the message and performs the LDAP acceptance validation within the SMTP conversation or the work queue. If the recipient is not found in the LDAP directory, you can configure the system to perform a delayed bounce or drop the message entirely.

For more information, see the "LDAP Queries" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide.*

SMTP Call-Ahead Recipient Validation

When you configure your Email Security appliance for SMTP call-ahead recipient validation, the Email Security appliance suspends the SMTP conversation with the sending MTA while it "calls ahead" to the SMTP server to verify the recipient. When the Cisco IronPort appliance queries the SMTP server, it returns the SMTP server's response to the Email Security appliance. The Email Security appliance resumes the SMTP conversation and sends a response to the sending MTA, allowing the conversation to continue or dropping the connection based on the SMTP server response (and settings you configure in the SMTP Call-Ahead profile).

For more information, see the "Validating Recipients Using an SMTP Server" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Work Queue / Routing

The Work Queue is where the received message is processed before moving to the delivery phase. Processing includes masquerading, routing, filtering, safelist/blocklist scanning, anti-spam and anti-virus scanning, Outbreak Filters, and quarantining.

Note Data loss prevention (DLP) scanning is only available for outgoing messages. For information on where DLP message scanning occurs in the Work Queue, see Message Splintering, page 6-4.

Email Pipeline and Security Services

Note, as a general rule, changes to security services (anti-spam scanning, anti-virus scanning, and Outbreak Filters) do not affect messages already in the work queue. As an example:

If a message bypasses anti-virus scanning when it first enters the pipeline because of any of these reasons:

- anti-virus scanning was not enabled globally for the appliance, or
- the HAT policy was to skip anti-virus scanning, or
- there was a message filter that caused the message to bypass anti-virus scanning,

then the message will not be anti-virus scanned upon release from the quarantine, regardless of whether anti-virus scanning has been re-enabled. However, messages that bypass anti-virus scanning due to mail policies may be anti-virus scanned upon release from a quarantine, as the mail policy's settings may have changed while the message was in the quarantine. For example, if a message bypasses anti-virus scanning due to a mail policy and is quarantined, then, prior to release from the quarantine, the mail policy is updated to include anti-virus scanning, the message will be anti-virus scanned upon release from the quarantine.

Similarly, suppose you had inadvertently disabled anti-spam scanning globally (or within the HAT), and you notice this after mail is in the work queue. Enabling anti-spam at that point will not cause the messages in the work queue to be anti-spam scanned.

LDAP Recipient Acceptance

You can use your existing LDAP infrastructure to define how the recipient email address of incoming messages (on a public listener) should be handled during the SMTP conversation or within the workqueue. See "Accept Queries" in the "Customizing Listeners" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*. This allows the Cisco IronPort appliance to combat directory harvest attacks (DHAP) in a unique way: the system accepts the message and performs the LDAP acceptance validation within the SMTP conversation or the work queue. If the recipient is not found in the LDAP directory, you can configure the system to perform a delayed bounce or drop the message entirely.

For more information, see the "LDAP Queries" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Masquerading or LDAP Masquerading

Masquerading is a feature that rewrites the envelope sender (also known as the sender, or `MAIL FROM`) and the To:, From:, and/or CC: headers on email processed by a private listener according to a table you construct. You can specify different masquerading parameters for each listener you create in one of two ways: via a static mapping table, or via an LDAP query.

For more information about masquerading via a static mapping table, see "Configuring Masquerading" in the "Configuring Routing and Delivery Features" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

For more information about masquerading via an LDAP query, see the "LDAP Queries" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

LDAP Routing

You can configure your Cisco IronPort appliance to route messages to the appropriate address and/or mail host based upon the information available in LDAP directories on your network.

For more information, see "LDAP Queries" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Message Filters

Message filters allow you to create special rules describing how to handle messages and attachments as they are received. Filter rules identify messages based on message or attachment content, information about the network, message envelope, message headers, or message body. Filter actions allow messages to be dropped, bounced, archived, quarantined, blind carbon copied, or altered.

For more information, see the "Using Message Filters to Enforce Email Policies" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Multi-recipient messages are "splintered" after this phase, prior to Email Security Manager. Splintering messages refers to creating splinter copies of emails with single recipients, for processing via Email Security Manager.

Email Security Manager (Per-Recipient Scanning)

Safelist/Blocklist Scanning

End user safelists and blocklists are created by end users and stored in a database that is checked prior to anti-spam scanning. Each end user can identify domains, sub domains or email addresses that they wish to always treat as spam or never treat as spam. If a sender address is part of an end users safelist, anti-spam scanning is skipped, and if the sender address is listed in the blocklist, the message may be quarantined or dropped depending on administrator settings. For more information about configuring safelists and blocklists, see the "Quarantines" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.

Anti-Spam

The Anti-Spam feature involves Cisco IronPort Anti-Spam scanning. Anti-spam scanning offers complete, Internet-wide, server-side anti-spam protection. It actively identifies and defuses spam attacks before they inconvenience your users and overwhelm or damage your network, allowing you to remove unwanted mail before it reaches your users' inboxes, without violating their privacy.

Anti-spam scanning can be configured to deliver mail to the Cisco IronPort Spam Quarantine (either on- or off-box). Messages released from the Cisco IronPort Spam Quarantine proceed directly to the destination queue, skipping any further work queue processing in the email pipeline.

See Chapter 9, "Anti-Spam" for more information.

Anti-Virus

Your Cisco IronPort appliance includes integrated virus scanning engines. You can configure the appliance to scan messages and attachments for viruses on a per-"mail policy" basis. You can configure the appliance to do the following when a virus is found:

- attempt to repair the attachment
- drop the attachment
- modify the subject header
- add an additional X- header
- send the message to a different address or mailhost

- archive the message
- delete the message

Messages released from quarantines (see Quarantines, page 4-9) are scanned for viruses. See Chapter 8, "Anti-Virus" for more information about Anti-Virus scanning.

Content Filters

You can create content filters to be applied to messages on a per-recipient or per-sender basis. Content filters are similar to message filters, except that they are applied later in the email pipeline — after a message has been "splintered" into a number of separate messages for each matching Email Security Manager policy. The functionality of content filters is applied after message filters processing and anti-spam and anti-virus scanning have been performed on a message.

For more information about Content Filters, see Content Filters Overview, page 6-6.

Outbreak Filters

Cisco IronPort's Outbreak Filters feature includes special filters that act proactively to provide a critical first layer of defense against new outbreaks. Based on Outbreak Rules published by Cisco IronPort, messages with attachments of specific filetypes can be sent to a quarantine named Outbreak.

Messages in the Outbreak quarantine are processed like any other message in a quarantine. For more information about quarantines and the Work Queue, see Quarantines, page 4-9.

See Chapter 10, "Outbreak Filters" for more information.

Quarantines

Cisco IronPort AsyncOS allows you to filter incoming or outgoing messages and place them into quarantines. Quarantines are special queues or repositories used to hold and process messages. Messages in quarantines can be delivered or deleted, based on how you configure the quarantine.

The following Work Queue features can send messages to quarantines:

- Message Filters
- Anti-Virus
- Outbreak Filters
- Content Filters

Messages released from quarantines are re-scanned for viruses.

See the "Quarantines" chapter of the *Cisco IronPort AsyncOS for Email Daily Management Guide* for more information.

Delivery

The delivery phase of the Email Pipeline focuses on the final phase of email processing, including limiting connections, bounces, and recipients.

Virtual gateways

The Cisco IronPort Virtual Gateway technology enables users to separate the Cisco IronPort appliance into multiple Virtual Gateway addresses from which to send and receive email. Each Virtual Gateway address is given a distinct IP address, hostname and domain, and email delivery queue.

For more information, see "Using Virtual Gateway™ Technology" in the "Configuring Routing and Delivery Features" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Delivery Limits

Use the `deliveryconfig` command to set limits on delivery, based on which IP interface to use when delivering and the maximum number of concurrent connections the appliance makes for outbound message delivery.

For more information, see "Set Email Delivery Parameters" in the "Configuring Routing and Delivery Features" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Domain-Based Limits

For each domain, you can assign a maximum number of connections and recipients that will never be exceeded by the system in a given time period. This "good neighbor" table is defined through the Mail Policies > Destination Controls page (or the `destconfig` command).

For more information, see "Controlling Email Delivery" in the "Configuring Routing and Delivery Features" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide.*

Domain-Based Routing

Use the Network > SMTP Routes page (or the `smtproutes` command) to redirect all email for a particular domain to a specific mail exchange (MX) host, without rewriting the envelope recipient.

For more information, see "Routing Email for Local Domains" in the "Configuring Routing and Delivery Features" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Global Unsubscribe

Use Global Unsubscribe to ensure that specific recipients, recipient domains, or IP addresses never receive messages from the Cisco IronPort appliance. If Global Unsubscribe is enabled, the system will check all recipient addresses against a list of "globally unsubscribed" users, domains, email addresses, and IP Addresses. Matching emails are not sent.

For more information, see "Using Global Unsubscribe" in the "Configuring Routing and Delivery Features" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Bounce Limits

You use the Network > Bounce Profiles page (or the `bounceconfig` command) to configure how Cisco IronPort AsyncOS handles hard and soft conversational bounces for each listener you create. You create bounce profiles and then apply profiles to each listener using the Network > Listeners page (or the `listenerconfig` command). You can also assign bounce profiles to specific messages using message filters.

For more information about bounce profiles, see "Directing Bounced Email" in the "Configuring Routing and Delivery Features" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

CHAPTER 5

Configuring the Gateway to Receive Email

After you have configured the basic settings of your Cisco IronPort appliance via the GUI's System Setup Wizard (or the CLI `systemsetup` command), you are now ready to begin tailoring the configuration of your Cisco IronPort Email Security appliance to receive email. This chapter discusses, in detail, all of the features available to you as you begin to configure listeners on the appliance to handle receiving email.

The concept of the Host Access Table (HAT) is introduced. The Host Access Tables (HATs) of public listeners — with their specific sender groups and mail flow policies — provide the underlying framework that makes possible the Mail Flow Monitor feature. ("Using Email Security Monitor" in the *Cisco IronPort AsyncOS for Email Daily Management Guide* describes the Mail Flow Monitor feature in detail.)

Receiving Email with Listeners

The Cisco IronPort AsyncOS operating system allows the Cisco IronPort appliance to function as the inbound email gateway for your enterprise, servicing SMTP connections from the Internet, accepting messages, and relaying messages to the appropriate systems.

In this configuration, you enable *listeners* to service these connections. A listener describes an email processing service that will be configured on a particular IP interface. Listeners only apply to email entering the Cisco IronPort appliance — either from the internal systems within your network or from the Internet. Cisco IronPort 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 injector" or even a "SMTP daemon" running on a specific port for each IP address you specify (including the initial addresses you configured with the `systemsetup` command).

Mail delivery policies cannot be configured so that mail is delivered to multiple ports on a single IP address (for example, port 25 for normal delivery and port 6025 for Cisco IronPort Spam quarantine). Cisco recommends running each delivery option on a separate IP address or host. Further, it is not possible to use the same hostname for regular email delivery and quarantine delivery.

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.

The System Setup Wizard or the `systemsetup` command (CLI) initially configures the *IP interfaces* that run on the available *Ethernet interfaces* on the Cisco IronPort appliance. On Cisco IronPort C150 and C160 appliances, these Ethernet interfaces are labeled Data1 and Data2. On all other Cisco IronPort appliances, they are labeled Data1, Data2, and Management. You can edit these interfaces at a later time via the IP Interfaces page on the Network menu or the `interfaceconfig` command. If you have completed the GUI's System Setup Wizard (or the `systemsetup` command) and committed the changes, at least one listener should already be configured on the appliance. (Refer to the settings you entered in the Step 3: Network, page 3-17.) The specific addresses to accept mail for were entered at that time, as well as the first SMTP Routes (Network > SMTP Routes or `smtproutes)` entry.

Note When you create a new listener via the System Setup Wizard, AsyncOS creates the listener with default values. However, when you create a listener manually, AsyncOS does not use these default SBRS values.

Use the Listeners page (Network > Listeners) or the `listenerconfig` command to configure listeners that run over available IP interfaces on the Cisco IronPort appliance. For more information about creating and configuring listeners, see the "Customizing Listeners" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*. In "Using Virtual Gateway™ Technology" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*, the Cisco IronPort Virtual Gateway technology is explained, in which you can further define and group IP interfaces over one or many IP addresses, or groups of IP addresses.

Figure 5-1 Relationship Between Listeners, IP Interfaces, and Physical Ethernet Interfaces

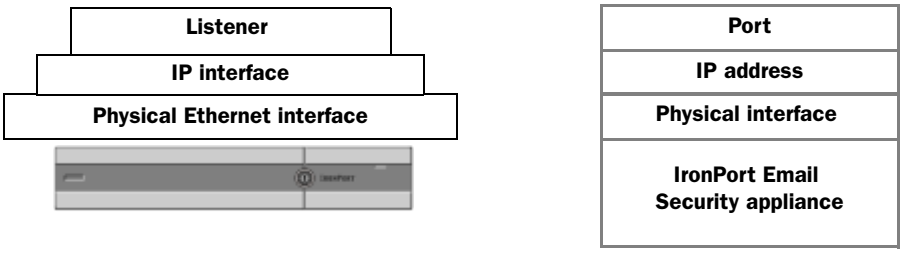

Enterprise Gateway Configuration

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 5-2 Using the Cisco IronPort Appliance as an Enterprise Gateway

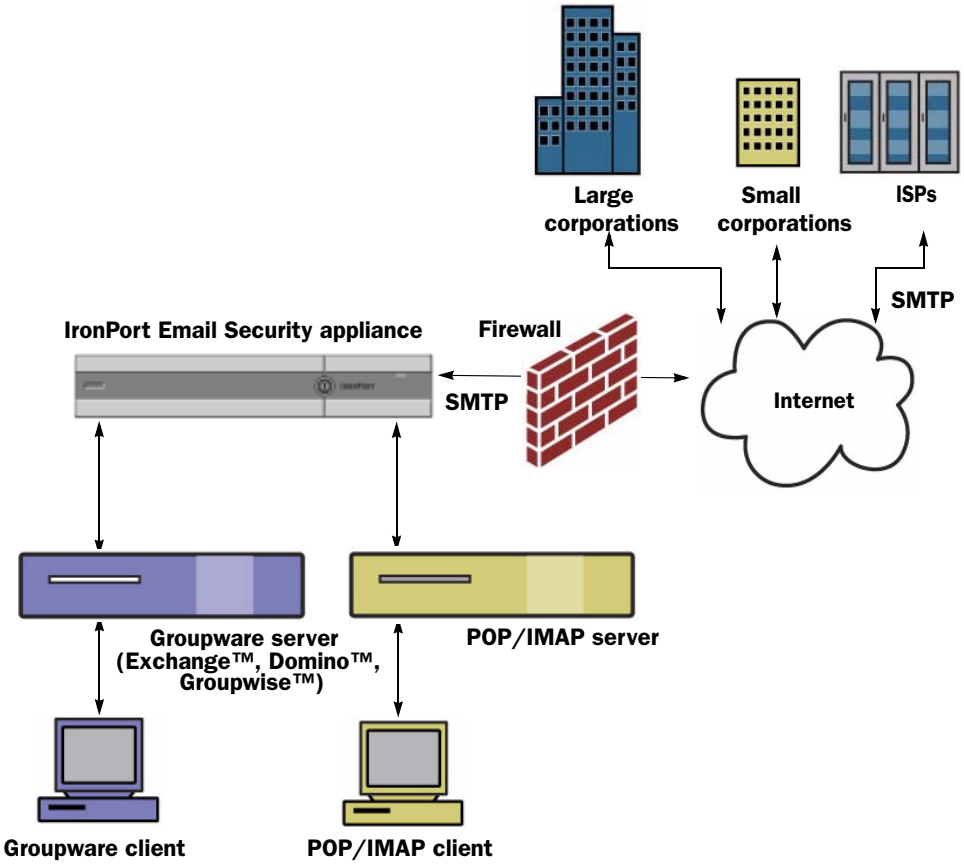

In this configuration, at least two listeners are required:

- 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)

Public and Private Listeners

Consider the first listener a "public" listener and the second listener a "private" listener. Cisco IronPort AsyncOS differentiates between public listeners — which by default have the characteristics for receiving email from the Internet — and private listeners that are intended to accept email only from internal (groupware, POP/IMAP, and other message generation) systems. Public and private listeners, by default, have different features available to them and different default settings. 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. The same illustration, with listeners, is shown in Figure 3-3.

Figure 5-3 Public and Private Listeners for an Enterprise Gateway

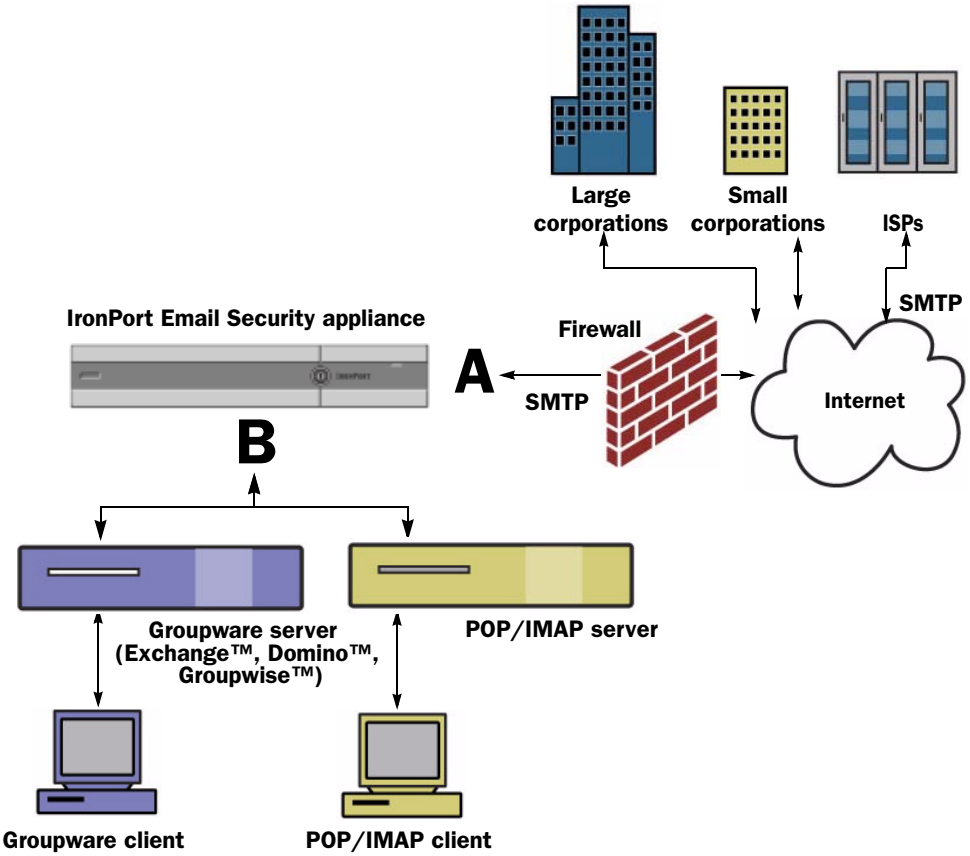

In Figure 3-3, one public listener (A) and one private listener (B) are configured on the appliance in this Enterprise Gateway configuration.

Figure 3-4 further illustrates the differences between the default settings of public and private listeners. A public listener is intended to receive email from the internet. The public listener receives connections from *many* hosts and directs messages to a *limited* number of recipients. Conversely, a private listener is intended to receive email from your internal network. The private listener receives connections from a limited (known) number of hosts and directs messages to *many* recipients.

C10/100 customers: 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.

Later in the chapter, these differences will be demonstrated in the Host Access Table and Recipient Access Table for each type of listener.

Figure 5-4 Public and Private Listeners

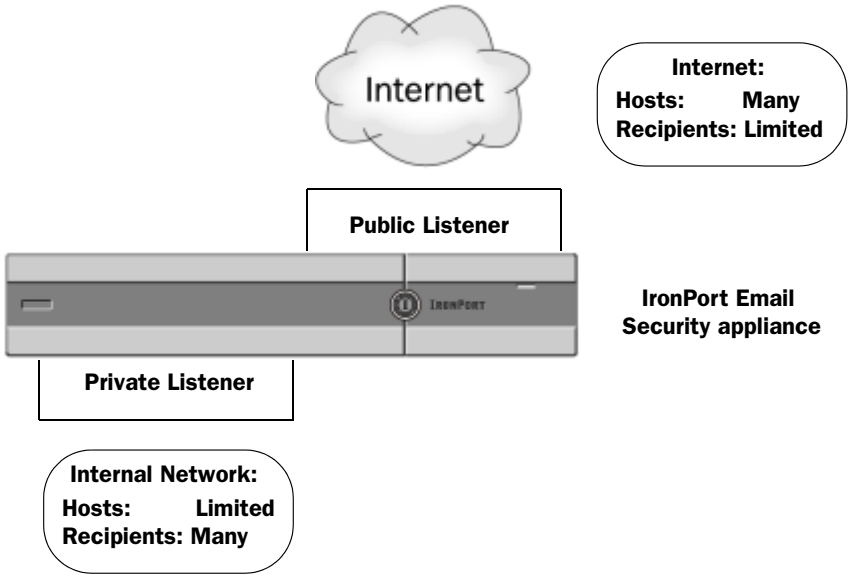

Figure 5-5 illustrates a typical Email Gateway configuration created by the System Setup Wizard Setup Wizard (or CLI `systemsetup` command) on Cisco IronPort X1050/1060/1070, C650/660/670, and C350/360/370 appliances. Two listeners are created: a public listener to serve inbound connections on one interface and a private listener to serve outbound connections on a second IP interface.

Figure 5-6 illustrates a typical Email Gateway configuration created by the System Setup Wizard (or CLI `systemsetup` command) on an Cisco IronPort C150/160 appliance. One public listener on a single IP interface is created to serve both inbound and outbound connections.

Figure 5-5 Public and Private Listeners on X1050/1060/1070, C650/660/670, C350/360/370 Appliances

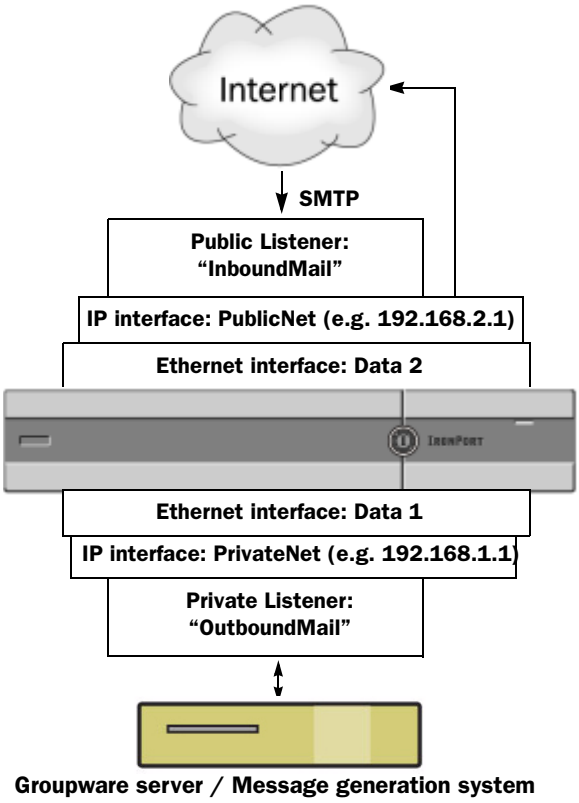

Note This public listener uses SMTP protocol on Port 25 of the PublicNet IP interface on the Data2 Ethernet interface to accept messages from the Internet.

IP interface PublicNet sends messages to destination hosts on the Internet.

IP interface PrivateNet sends messages to internal mail hosts.

Note This private listener uses SMTP protocol on Port 25 of the PrivateNet IP interface on the Data1 Ethernet interface to accept messages from internal systems in the `.example.com` domain.

Figure 5-6 Public Listener on C150/160 Appliance

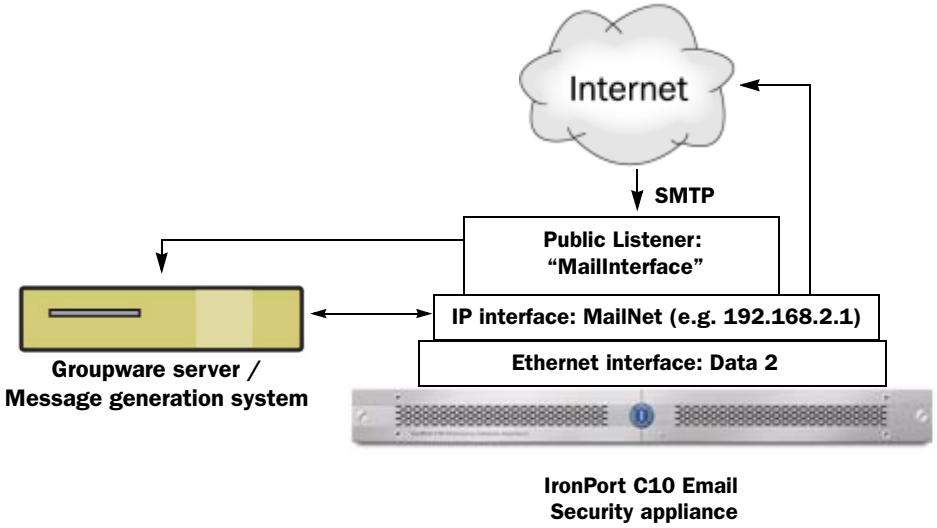

Note This public listener uses SMTP protocol on Port 25 of the PublicNet IP interface on the Data2 Ethernet interface to accept messages from the Internet and to relay messages from internal systems in the `.example.com` domain.
IP interface MailNet sends messages to destination hosts on the Internet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

The Host Access Table (HAT): Sender Groups and Mail Flow Policies

Each listener that is configured on an appliance has properties that you can configure to modify the behavior of the message it receives. As discussed in the Overview: Email Pipeline, page 4-1, one of the first configurable features that influences a listener's behavior is its Host Access Table (HAT).

The HAT maintains a set of rules that control incoming connections from remote hosts for a listener. Every listener you create has its own HAT. HATs are defined for public and private listeners.

Entries in HAT are defined by this basic syntax:

Table 5-1 Basic HAT Syntax

Remote Host Definition	Rule

The *remote host definition* is the way in which a remote host that is attempting to connect to the listener is defined (for example, by a single IP address).

A *rule* defines whether the remote host specified can or cannot connect to the listener.

Extending the basic syntax, HATs in AsyncOS support the ability to create named sets of remote host definitions; these are called *sender groups*. Named sets of access rules combined with parameter sets are called *mail flow policies*. This extended syntax is illustrated in Table 5-2:

Table 5-2 Advanced HAT Syntax

Sender Group:	Mail Flow Policy:
Remote Host	Access Rule + Parameters
Remote Host	
Remote Host	
...	

The order that rules appear in the HAT is important. The HAT is read from top to bottom for each host that attempts to connect to the listener. If a rule matches a connecting host, the action is taken for that connection immediately.

Predefined and custom entries you place in the HAT are entered above the final "ALL" host entry.

Default HAT Entries

For all public listeners you create, by default, the HAT is set to *accept* email from *all* hosts. For all private listeners you create, by default, the HAT is set up to relay email from the host(s) you specify, and reject *all* other hosts.

Note By rejecting all hosts other than the ones you specify, the `listenerconfig` and `systemsetup` commands prevent you from unintentionally configuring your system as an "open relay." An open relay (sometimes called an "insecure relay" or a "third party" relay) is an SMTP email server that allows third-party relay of email messages. By processing email that is neither for nor from a local user, an open relay makes it possible for an unscrupulous sender to route large volumes of spam through your gateway.

Mail Flow Policies: Access Rules and Parameters

Mail Flow Policies of the HAT allow you to control or limit the rates at which the listener will receive mail from remote hosts. You can also modify the SMTP codes and responses communicated during the SMTP conversation.

The HAT has four basic access rules for acting on connections from remote hosts:

Step 1 `ACCEPT`

Connection is accepted, and email acceptance is then further restricted by listener settings, including the Recipient Access Table (for public listeners).

Step 2 `REJECT`

Connection is initially accepted, but the client attempting to connect gets a 4XX or 5XX greeting. No email is accepted.

Note You can also configure AsyncOS to perform this 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. This setting is configured from the CLI `listenerconfig --> setup` command. For more information, see "Customizing Listeners" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Step 3 `TCPREFUSE`

Connection is refused at the TCP level.

Step 4 `RELAY`

Connection is accepted. Receiving for any recipient is allowed and is not constrained by the Recipient Access Table.

- `CONTINUE`

  The mapping in the HAT is ignored, and processing of the HAT continues. If the incoming connection matches a later entry that is not CONTINUE, that entry is used instead. The CONTINUE rule is used to facilitate the editing of the HAT in the Graphical User Interface (GUI). For more information, see Adding a New Sender Group, page 5-31.

In addition to these basic access control parameters, the following parameters are available for listeners you create. Parameters combined with an access rule (ACCEPT or REJECT) are called *mail flow policies*. A mail flow policy is a way of expressing a group of HAT parameters (access rule, followed by connection parameters, rate limiting parameters, custom SMTP codes and responses, and anti-spam, anti-virus, encryption, and authentication parameters).

Mail flow policies are then mapped to sender groups as entries in a listener's HAT.

Table 5-3 HAT Mail Flow Policy Parameters

Parameter	Description
Connections	
Maximum message size	The maximum size of a message that will be accepted by this listener. The smallest possible maximum message size is 1 kilobyte.
Maximum concurrent connections from a single IP	The maximum number of concurrent connections allowed to connect to this listener from a single IP address.
Maximum messages per connection	The maximum number of messages that can be sent through this listener per connection from a remote host.
Maximum recipients per message	That maximum number of recipients per message that will be accepted from this host.
SMTP Banner	
Custom SMTP Banner Code	The SMTP code returned when a connection is established with this listener.
Custom SMTP Banner Text	The SMTP banner text returned when a connection is established with this listener.
Custom SMTP Reject Banner Code	The SMTP code returned when a connection is rejected by this listener.
Custom SMTP Reject Banner Text	The SMTP banner text returned when a connection is rejected by this listener.

Table 5-3 HAT Mail Flow Policy Parameters (Continued)

Parameter	Description
Override SMTP Banner Host Name	By default, the appliance will include the hostname associated with the interface of the listener when displaying the SMTP banner to remote hosts (for example: `220-`*hostname* `ESMTP`). You may choose to override this banner by entering a different hostname here. Additionally, you may leave the hostname field blank to choose *not* to display a hostname in the banner.
Rate Limit for Hosts	
Max. Recipients per Hour	The maximum number of recipients per hour this listener will receive from a remote host. The number of recipients per sender IP address is tracked globally. Each listener tracks its own rate limiting threshold; however, because all listeners validate against a single counter, it is more likely that the rate limit will be exceeded if the same IP address (sender) is connecting to multiple listeners.
Max. Recipients per Hour Code	The SMTP code returned when a host exceeds the maximum number of recipients per hour defined for this listener.
Max. Recipients Per Hour Exceeded Text	The SMTP banner text returned when a host exceeds the maximum number of recipients per hour defined for this listener.
Rate Limit for Sender	
Max. Recipients per Time Interval	The maximum number of recipients during a specified time period that this listener will receive from a unique envelope sender, based on the mail-from address. The number of recipients is tracked globally. Each listener tracks its own rate limiting threshold; however, because all listeners validate against a single counter, it is more likely that the rate limit will be exceeded if messages from the same mail-from address are received by multiple listeners. Select whether to use the default maximum recipients, accept unlimited recipients, or specify another maximum number of recipients. Use the Default Mail Flow Policy settings to specify the maximum number of recipients and the time interval that will be used by the other mail flow policies by default. The time interval can only be specified using the Default Mail Flow Policy.
Sender Rate Limit Exceeded Error Code	The SMTP code returned when an envelope exceeds the maximum number of recipients for the time interval defined for this listener.
Sender Rate Limit Exceeded Error Text	The SMTP banner text returned when an envelope sender exceeds the maximum number of recipients for the time interval defined for this listener.
Exceptions	If you want certain envelope senders to be exempt from the defined rate limit, select an address list that contains the envelope senders. See Address Lists, page 5-39 for more information.
Flow Control	
Use SenderBase for Flow Control	Enable “look ups” to the Cisco IronPort SenderBase Reputation Service for this listener.

Table 5-3 HAT Mail Flow Policy Parameters (Continued)

Parameter	Description
Group by Similarity of IP Addresses: (significant bits 0-32)	Used to track and rate limit incoming mail on a per-IP address basis while managing entries in a listener's Host Access Table (HAT) in large CIDR blocks. You define a range of significant bits (from 0 to 32) by which to group similar IP addresses for the purposes of rate limiting, while still maintaining an individual counter for each IP address within that range. Requires "Use SenderBase" to be disabled. For more information about HAT significant bits, see "HAT Significant Bits Feature" in the "Configuring Routing and Delivery Features" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.
Directory Harvest Attack Prevention (DHAP)	
Directory Harvest Attack Prevention: Maximum Invalid Recipients Per Hour	The maximum number of invalid recipients per hour this listener will receive from a remote host. This threshold represents the total number of RAT rejections and SMTP call-ahead server rejections combined with the total number of messages to invalid LDAP recipients dropped in the SMTP conversation or bounced in the work queue (as configured in the LDAP accept settings on the associated listener). For more information on configuring DHAP for LDAP accept queries, see "LDAP Queries" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.
Directory Harvest Attack Prevention: Drop Connection if DHAP threshold is Reached within an SMTP Conversation	The Cisco IronPort appliance will drop a connection to a host if the threshold of invalid recipients is reached.
Max. Invalid Recipients Per Hour Code:	Specify the code to use when dropping connections. The default code is 550.
Max. Invalid Recipients Per Hour Text:	Specify the text to use for dropped connections. The default text is "Too many invalid recipients."
Drop Connection if DHAP threshold is reached within an SMTP Conversation	Enable to drop connections if the DHAP threshold is reached within an SMTP conversation.
Max. Invalid Recipients Per Hour Code	Specify the code to use when dropping connections due to DHAP within an SMTP conversation. The default code is 550.
Max. Invalid Recipients Per Hour Text:	Specify the text to use when dropping connections due to DHAP within an SMTP conversation.
Spam Detection	
Anti-spam scanning	Enable anti-spam scanning on this listener.
Virus Detection	
Anti-virus scanning	Enable the anti-virus scanning on this listener.
Encryption and Authentication	

Table 5-3 HAT Mail Flow Policy Parameters (Continued)

Parameter	Description
Allow TLS Connections	Deny, Prefer, or Require Transport Layer Security (TLS) in SMTP conversations for this listener. If you select Preferred, you can make TLS mandatory for envelope senders from a specific domain or with a specific email address by selecting an Address List that specifies those domains and email addresses. When an envelope sender matching a domain or address in this list tries to send a message over a connection that does not use TLS, the appliance rejects the connection and the sender will have to try again using TLS. For information on creating an address list, see Address Lists, page 5-39.
SMTP Authentication	Allows, disallow, or requires SMTP Authentication from remote hosts connecting to the listener. SMTP Authentication is described in detail in the "LDAP Queries" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.
If Both TLS and SMTP Authentication are enabled:	Require TLS to offer SMTP Authentication.
Domain Key Signing	
Domain Key/ DKIM Signing	Enable Domain Keys or DKIM signing on this listener (ACCEPT and RELAY only).
DKIM Verification	Enable DKIM verification.
SPF/SIDF Verification	
Enable SPF/SIDF Verification	Enable SPF/SIDF signing on this listener. For more information, see the "Email Authentication" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.
Conformance Level	Set the SPF/SIDF conformance level. You can choose from SPF, SIDF or SIDF Compatible. For details, see the "Email Authentication" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.
Downgrade PRA verification result if 'Resent-Sender:' or 'Resent-From:' were used:	If you choose a conformance level of SIDF compatible, configure whether you want to downgrade Pass result of the PRA Identity verification to None if there are Resent-Sender: or Resent-From: headers present in the message. You may choose this option for security purposes.
HELO Test	Configure whether you want to perform a test against the HELO identity (Use this for SPF and SIDF Compatible conformance levels).
Untagged Bounces	
Consider Untagged Bounces to be Valid	Applies only if bounce verification tagging (discussed in the "Configuring Routing and Delivery Features" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*) is enabled. By default, the appliance considers untagged bounces invalid and either rejects the bounce or adds a custom header, depending on the Bounce Verification settings. If you choose to consider untagged bounces to be valid, the appliance accepts the bounce message.

Table 5-3 HAT Mail Flow Policy Parameters (Continued)

Parameter	Description
Envelope Sender DNS Verification	
	See Sender Verification, page 5-40.
Exception Table	
Use Exception Table	Use the sender verification domain exception table. You can only have one exception table, but you can enable it per mail flow policy. See Sender Verification Exception Table, page 5-43 for more information.

By default, these parameters are set to the following default values shown in Table 5-5 and Table 5-6 for each listener on the appliance.

Note If anti-spam or anti-virus scanning is enabled globally in the HAT, messages are flagged for anti-spam or anti-virus scanning as they are accepted by the Cisco IronPort appliance. If anti-spam or anti-virus scanning is disabled after the message is accepted, the message will still be subject to scanning when it leaves the work queue.

HAT Variable Syntax

Table 5-4 defines a set of variables that can also be used in conjunction with the custom SMTP and Rate Limiting banners defined for a Mail Flow Policy. Variable names are case-insensitive. (That is, `$group` is equivalent to `$Group`.)

Table 5-4 HAT Variable Syntax

Variable	Definition
`$Group`	Replaced by the name of the sender group that was matched in the HAT. If the sender group has no name, “None” is displayed.
`$Hostname`	Replaced by the remote hostname if and only if is has been validated by the Cisco IronPort appliance. If the reverse DNS lookup of the IP address is successful but returns no hostname, then “None” is displayed. If the reverse DNS lookup fails (for example, if the DNS server cannot be reached, or no DNS server has been configured) then “Unknown” is displayed.
`$OrgID`	Replaced by the SenderBase Organization ID (an integer value). If the Cisco IronPort appliance cannot obtain a SenderBase Organization ID, or if the SenderBase Reputation Service did not return a value, “None” is displayed.
`$RemoteIP`	Replaced by the IP address of the remote client.
`$HATEntry`	Replaced by the entry in the HAT that the remote client matched.

Using HAT Variables

Note These variables can be used with the `smtp_banner_text` and `max_rcpts_per_hour_text` advanced HAT parameters shown in Table 1-3 of the “Customizing Listeners” chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Using these variables, you could edit the custom SMTP banner response text for accepted connections in the $TRUSTED policy in the GUI:

Figure 5-7 Using HAT Variables

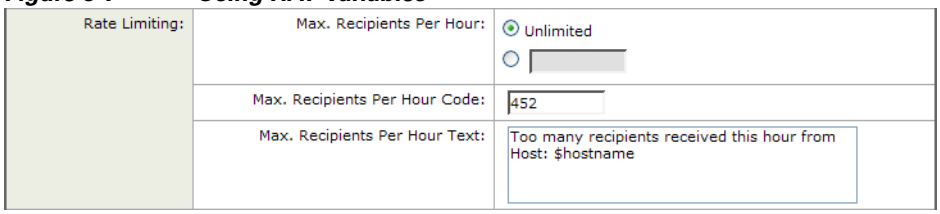

Or like this, in the CLI:

```
Would you like to specify a custom SMTP response?  [Y]> y


Enter the SMTP code to use in the response. 220 is the standard code.

[220]> 200


Enter your custom SMTP response.  Press Enter on a blank line to finish.

You've connected from the hostname: $Hostname, IP address of: $RemoteIP, matched the
group: $Group, $HATEntry and the SenderBase Organization: $OrgID.
```

Testing HAT Variables

To test these variables, add the IP address of a known, trusted machine to the $WHITELIST sender group of a listener on the Cisco IronPort appliance. Then, connect from that machine via telnet. You can see the variable substitution in the SMTP response. For example:

```
# telnet IP_address_of_IronPort_Appliance

220 hostname ESMTP

200 You've connected from the hostname: hostname, IP address of:
IP-address_of_connecting_machine, matched the group: WHITELIST, 10.1.1.1 the SenderBase
Organization: OrgID.
```

Viewing Default Mail Flow Policies

Figure 5-8 shows the default policy parameters for a public listener.

Step 1 Access the GUI (see Accessing the GUI, page 2-2).

Step 2 Click Mail Policies > Mail Flow Policies.

The Mail Flow Policies page is displayed. If listeners are configured, the mail flow policies defined for the first alphabetical listener are displayed.

Figure 5-8 Mail Flow Policies Page

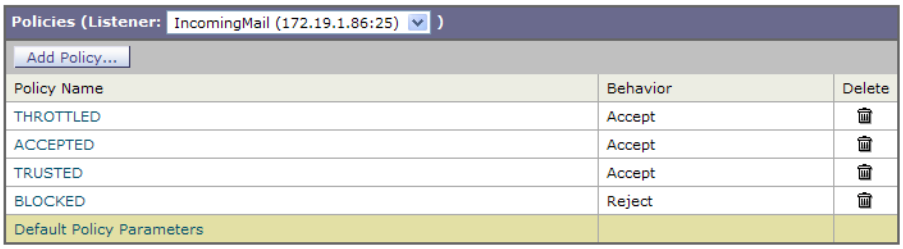

Step 3 Click the Default Policy Parameters link.

The default policy parameters page is displayed. See Figure 5-9.

Figure 5-9 Default Policy Parameters for a Public Listener (one of two)

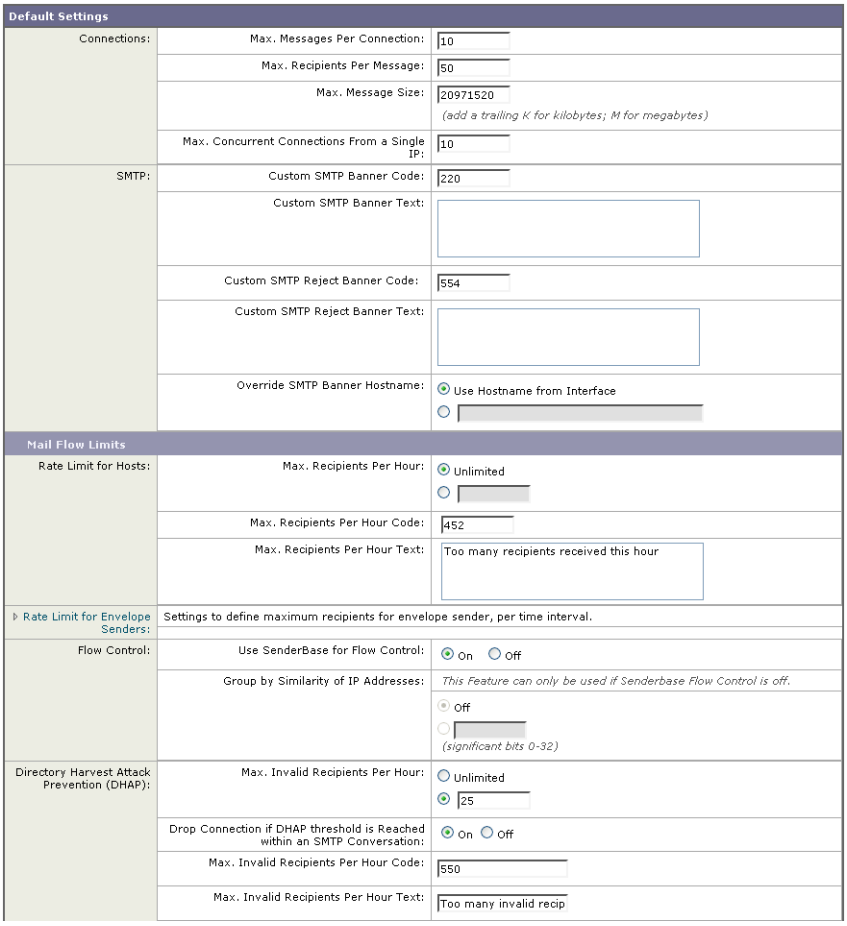

Figure 5-10 Default Policy Parameters for a Public Listener (two of two)

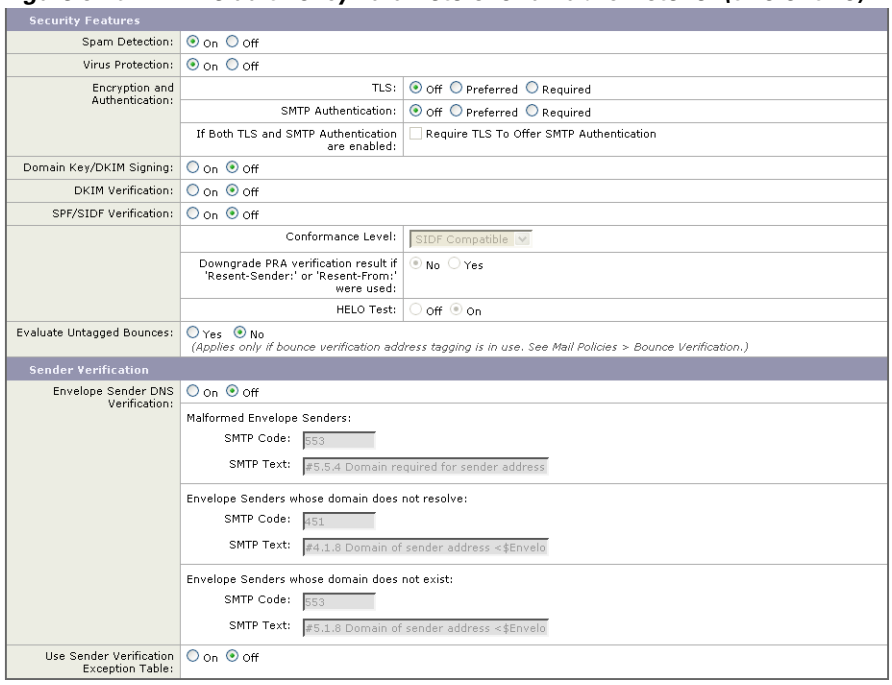

Default Policy Parameters for Listeners

The following table lists the default parameters for public listeners.

Table 5-5 HAT Default Policy Parameters for Public Listeners

Parameters	Default Value
Maximum message size:	20 MB
Max. concurrent connections allowed to this listener:	10 connections
Maximum messages per connection:	10 messages
Maximum recipients per message:	50 recipients
SMTP Banner Code:	220
SMTP Banner Text:	"*hostname* `ESMTP`"
SMTP Reject Banner Code:	554
SMTP Reject Banner Text:	"`Access Denied`"
Override SMTP Banner Hostname	Use hostname from Interface
Maximum Recipients per Hour:	No default. User-defined.
Maximum Recipients per Hour Code:	452
Maximum Recipients per Hour Text:	"`Too many recipients received this hour`"
Maximum Recipients Per Time Interval:	Unlimited
Maximum Sender Rate Limit Error Code	452

Table 5-5 HAT Default Policy Parameters for Public Listeners

Parameters	Default Value
Sender Rate Limit Error Text	"`Too many recipients received from the sender`"
Exceptions	None
Directory Harvest Attack Prevention	OFF
Use SenderBase:	ON
Group by Similarity of IP address:	DISABLED
Use anti-spam scanning:	ON (If anti-spam enabled)
Use anti-virus scanning:	ON (If anti-virus enabled)
Allow TLS Connections:	NO
Override Hostname	NO
SMTP Auth	OFF
Domainkey/DKIM Signing	OFF
DKIM Verification	OFF
SPF/SIDF Verification	OFF
Envelope Sender DNS Verification	OFF
Use Exception Table	OFF

The following table lists the default parameters for private listeners.

Table 5-6 HAT Default Policy Parameters for Private Listeners

Parameters	Default Value
Maximum messages per connection:	10,000 messages
Maximum recipients per message:	100,000 recipients
Maximum message size:	100 MB (104857600 bytes)
Max. concurrent connections from a single IP	50 connections
SMTP Banner Code:	220
SMTP Banner Text:	"*hostname* `ESMTP`"
SMTP Reject Banner Code:	554
SMTP Reject Banner Text:	"`Access Denied`"
Override SMTP Banner Hostname	Use Hostname from Interface
Use SenderBase:	OFF
Maximum Recipients per Hour:	No default. User-defined.
Maximum Recipients per Hour Code:	452
Maximum Recipients per Hour Text:	"`Too many recipients received this hour`"
Maximum Recipients Per Time Interval:	Unlimited

Table 5-6 HAT Default Policy Parameters for Private Listeners

Parameters	Default Value
Maximum Sender Rate Limit Error Code	452
Sender Rate Limit Error Text	"Too many recipients received from the sender"
Exceptions	None
Group by Similarity of IP address:	OFF
Directory Harvest Attack Prevention	OFF
Use anti-spam scanning:	OFF (If anti-spam enabled)
Use anti-virus scanning:	ON (If anti-virus enabled)
Allow TLS Connections:	NO
Override Hostname	NO
SMTP Auth	OFF
Domainkeys/DKIM Signing	OFF
DKIM Verification	OFF
SPF/SIDF Verification	OFF
Accept Untagged Bounces	NO
Envelope Sender DNS Verification	OFF
Use Exception Table	OFF

Sender Groups

HAT parameters are combined with an access rule to create a mail flow policy (see Figure 5-6Mail Flow Policies: Access Rules and Parameters, page 5-8). When you group together different HAT parameters and assign a name to them, you are defining a mail flow policy that can be applied to groups of senders.

A *sender group* is simply a list of senders gathered together for the purposes of handling email from those senders in the same way (that is, applying a mail flow policy to a group of senders). A sender group is a list of senders identified by:

- IP address (IPv4 or IPv6)
- IP range
- Specific host or domain name
- SenderBase Reputation Service "organization" classification
- SenderBase Reputation Score (SBRS) range (or lack of score)
- DNS List query response

See Table 5-7 for the syntax of defining remote hosts (sender entries) that make up sender groups. These sender entries are separated by commas in a listener's HAT. You assign a name for sender groups, as well as mail flow policies.

Together, sender groups and mail flow policies are defined in a listener's HAT. By default, your Cisco IronPort appliance ships with the predefined mail flow policies and sender groups described in Accessing Predefined Mail Flow Policies for Public Listeners, page 5-25.

In Chapter 6, "Email Security Manager" you will use the predefined sender groups and mail flow policies to quickly and powerfully classify the mail flowing through your gateway, enabling real-time changes to a listener's HAT.

Note The system acquires and verifies the validity of the remote host's IP address by performing a double DNS lookup. This consists of a reverse DNS (PTR) lookup on the IP address of the connecting host, followed by a forward DNS (A) lookup on the results of the PTR lookup. The system then checks that the results of the A lookup match the results of the PTR lookup. If the results do not match, or if an A record does not exist, the system only uses the IP address to match entries in the HAT.

Sender Group Syntax

Table 5-7 Defining Remote Hosts in the HAT: Sender Group Syntax

Syntax	Meaning
n:n:n:n:n:n:n:n	IPv6 address; does not need to include leading zeroes.
n:n:n:n:n:n:n:n- n:n:n:n:n:n:n:n n:n:n-n:n:n:n:n:n	Range of IPv6 addresses; does not need to include leading zeroes.
`n.n.n.n`	Full (complete) IPv4 Address

`n.n.n.`
 	
`n.n.n`
 	
`n.n.`
 	
`n.n`
 	
`n.`
 	
`n`
	Partial IPv4 address

`n.n.n.n-n`
 	
`n.n.n-n.`
 	
`n.n.n-n`
 	
`n.n-n.`
 	
`n.n-n`
 	
`n-n.`
 	
`n-n`
	Range of IPv4 addresses

`yourhost.example.com`	A fully-qualified domain name
`.partialhost`	Everything within the partialhost domain

`n/c`
 	
`n.n/c`
 	
`n.n.n/c`
 	
`n.n.n.n/c`
	IPv4 CIDR address block

n:n:n:n:n:n:n:n/c	IPv6 CIDR address block; does not need to include leading zeroes
`SBRS[n:n]` `SBRS[none]`	SenderBase Reputation Score. For more information, see Sender Groups defined by SenderBase Reputation Scores, page 5-23.
`SBO:n`	SenderBase Network Owner Identification Number. For more information, see Sender Groups defined by SenderBase Reputation Scores, page 5-23.
dnslist[dnsserver.domain]	DNS List query. For more information, see Sender Groups Defined by Querying DNS Lists in the HAT, page 5-24.
`ALL`	Special keyword that matches ALL addresses. This applies only to the ALL sender group, and is always included (but not listed).

Sender Groups Defined by Network Owners, Domains, and IP Addresses

Since the SMTP protocol has no built-in method for authenticating senders of email, senders of unsolicited bulk email have been successful at employing a number of tactics for hiding their identity. Examples include spoofing the Envelope Sender address on a message, using a forged HELO address, or simply rotating through different domain names. This leaves many mail administrators asking themselves the fundamental question, "Who is sending me all of this email?" To answer this question, the SenderBase Reputation Service has developed a unique hierarchy for aggregating identity-based information based on the IP address of the connecting host — the one thing that is almost impossible for a sender to forge in a message.

An **IP Address** is defined as the IP address of the sending mail host. The Email Security appliance supports both Internet Protocol version 4 (IPv4) and version 6 (IPv6) addresses.

A **Domain** is defined as an entity that uses hostnames with a given second-level domain name (for example, yahoo.com), as determined by a reverse (PTR) lookup on the IP address.

A **Network Owner** is defined as an entity (usually a company) that controls a block of IP addresses, as determined based on IP address space assignments from global registries such as ARIN (the American Registry for Internet Numbers) and other sources.

An **Organization** is defined as an entity that most closely controls a particular group of mail gateways within a network owner's IP block, as determined by SenderBase. An Organization may be the same as the Network Owner, a division within that Network Owner, or a customer of that Network Owner.

Setting Policies Based on the HAT

Table 5-8 lists some examples of network owners and organizations.

Table 5-8 Example of Network Owners and Organizations

Example Type	Network Owner	Organization
Network Service Provider	Level 3 Communications	Macromedia Inc. AllOutDeals.com GreatOffers.com
Email Service Provider	GE	GE Appliances GE Capital GE Mortgage
Commercial Sender	The Motley Fool	The Motley Fool

As network owners can range dramatically in size, the appropriate entity to base your mail flow policy on is the organization. The SenderBase Reputation Service has a unique understanding of the source of the email down to the organization level, which the Cisco IronPort appliance leverages to automatically apply policies based on the organization. In the example above, if a user specified "Level 3 Communications" as a sender group in the Host Access Table (HAT), SenderBase will enforce policies based on the individual organizations controlled by that network owner.

For example, in Table 3-7 above, if a user enters a limit of 10 recipients per hour for Level 3, the Cisco IronPort appliance will allow up to 10 recipients per hour for Macromedia Inc., Alloutdeals.com *and* Greatoffers.com (a total of 30 recipients per hour for the Level 3 network owner). The advantage of this approach is that if one of these organizations begins spamming, the other organizations controlled by Level 3 will not be impacted. Contrast this to the example of "The Motley Fool" network owner. If a user sets rate limiting to 10 recipients per hour, the Motley Fool network owner will receive a total limit of 10 recipients per hour.

The Cisco IronPort Mail Flow Monitor feature is a way of defining the sender and providing you with monitoring tools to create mail flow policy decisions about the sender. To create mail flow policy decisions about a given sender, ask these questions:

Step 1 **Which IP addresses are controlled by this sender?**

The first piece of information that the Mail Flow Monitor feature uses to control the inbound email processing is the answer to this question. The answer is derived by querying the SenderBase Reputation Service. The SenderBase Reputation Service provides information about the relative size of the sender (either the SenderBase network owner or the SenderBase organization). Answering this question assumes the following:

- Larger organizations tend to control more IP addresses, and send more legitimate email.

Step 2 **Depending on its size, how should the overall number of connections be allotted for this sender?**

- Larger organizations tend to control more IP addresses, and send more legitimate email. Therefore, they should be allotted more connections to your appliance.
- The sources of high-volume email are often ISPs, NSPs, companies that manage outsourced email delivery, or sources of unsolicited bulk email. ISPs, NSPS, and companies that manage outsourced email delivery are examples of organizations that control many IP addresses, and should be allotted more connections to your appliance. Senders of unsolicited bulk email usually do not control many IP addresses; rather, they send large volumes of mail through a few number of IP addresses. They should be allotted fewer connections to your appliance.

The Mail Flow Monitor feature uses its differentiation between SenderBase network owners and SenderBase organizations to determine how to allot connections per sender, based on logic in SenderBase. See the "Using Email Security Monitor" chapter in *Cisco IronPort AsyncOS for Email Daily Management Guide* for more information on using the Mail Flow Monitor feature.

Sender Groups defined by SenderBase Reputation Scores

The Cisco IronPort appliance can query the Cisco IronPort SenderBase Reputation Service to determine a sender's reputation score (SBRS). The SBRS is a numeric value assigned to an IP address, domain, or organization based on information from the SenderBase Reputation Service. The scale of the score ranges from -10.0 to +10.0, as described in Table 5-9.

Table 5-9 Definition of the SenderBase Reputation Score

Score	Meaning
-10.0	Most likely to be a source of spam
0	Neutral, or not enough information to make a recommendation
+10.0	Most likely to be a trustworthy sender
none	No data available for this sender (typically a source of spam)

Using the SBRS, you configure the Cisco IronPort appliance to apply mail flow policies to senders based on their trustworthiness. For example, all senders with a score less than -7.5 could be rejected. This is most easily accomplished via the GUI; see Creating a Sender Group with SenderBase Reputation Scores,

page 5-34. However, if you are modifying an exported HAT in a text file, the syntax for including SenderBase Reputation Scores is described in Table 5-10. See "Customizing Listeners" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Table 5-10 Syntax for SenderBase Reputation Scores within a HAT

`SBRS[n:n]`	SenderBase Reputation Score. Senders are identified by querying the SenderBase Reputation Service, and the scores are defined between the ranges.
SBRS[none]	Specify no SBRS (very new domains may not have SenderBase Reputation Scores yet).

Note Network owners added to a HAT via the GUI use the syntax `SBO:n`, where *n* is the network owner's unique identification number in the SenderBase Reputation Service.

Use the Network > Listeners page or `listenerconfig -> setup` command in the CLI to enable a listener to query the SenderBase Reputation Service. You can also define the timeout value that the appliance should wait when querying the SenderBase Reputation Service. Then, you can configure different policies to use look ups to the SenderBase Reputation Service by using the values in the Mail Policies Pages in the GUI or the `listenerconfig -> edit -> hostaccess` commands in the CLI.

Note You can also create message filters to specify "thresholds" for SenderBase Reputation Scores to further act upon messages processed by the system. For more information, see "SenderBase Reputation Rule," "Bypass Anti-Spam System Action," and "Bypass Anti-Virus System Action" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Sender Groups Defined by Querying DNS Lists in the HAT

You also have the ability in a listener's HAT to define a sender group as matching a query to a specific DNS List sever. The query is performed via DNS at the time of the remote client's connection. The ability to query a remote list also exists currently as a message filter rule (see "DNS List Rule" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*), but only once the message content has been received in full.

This mechanism allows you to configure a sender within a group that queries a DNS List so that you can adjust your mail flow policies accordingly. For example, you could reject connections or limit the behavior of the connecting domain.

Note Some DNS Lists use variable responses (for example, "127.0.0.1" versus "127.0.0.2" versus "127.0.0.3") to indicate various facts about the IP address being queried against. If you use the message filter DNS List rule (see "DNS List Rule" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*), you can compare the result of the query against different values. However, specifying a DNS List server to be queried in the HAT only supports a Boolean operation for simplicity (that is, does the IP address appear in the list or not)

Note Be sure to include brackets in the query in the CLI. Brackets are not necessary when specifying a DNS List query in the GUI. Use the `dnslistconfig` command in the CLI to test a query, configure general settings for DNL queries, or flush the current DNS list cache.

Note that this mechanism can be used to identify "good" connections as well as "bad" connections. For example, a query to query.bondedsender.org will match on connecting hosts who have posted a financial bond with Cisco IronPort Systems' Bonded Sender™ program to ensure the integrity of their email campaign. You could modify the default WHITELIST sender group to query the Bonded Sender program's DNS servers (which lists these legitimate email senders who have willingly posted bonds) and adjust the mail flow policy accordingly.

Accessing Predefined Mail Flow Policies for Public Listeners

When combined with an access rule (ACCEPT or REJECT), the parameters listed in Table 5-3 on page 5-9 are predefined as the following four mail flow policies for each *public* listener you create:

- $ACCEPTED
- $BLOCKED
- $THROTTLED
- $TRUSTED

Step 1 Access the GUI. (See Accessing the GUI, page 2-2.)

Step 2 Click Mail Policies > HAT Overview.

The Overview page is displayed. If listeners are configured, the Host Access Table overview page defined for the first alphabetical listener is displayed. Select the desired listener from the Listener list.

Figure 5-11 Predefined Mail Flow Policies for Public Listeners

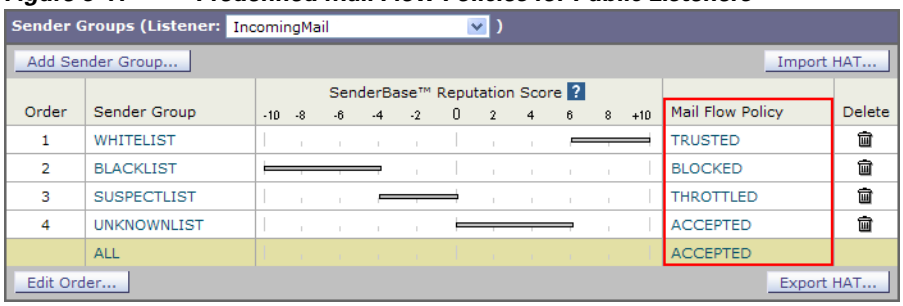

Step 3 Click the name of a Mail Flow Policy to view the connection behavior and parameters for that policy.

Note By default, C150/160 customers are prompted to create only one public listener during the `systemsetup` command. Public listeners created on Cisco IronPort C150/160 appliances also include a $RELAYED mail flow policy that is used to relay mail for internal systems (as shown in Figure 5-12). For more information, see RELAYLIST, page 5-30. The $RELAYLIST policy is shown only on private listeners on Cisco IronPort X1050/1060/1070, C650/660/670, and C350/360/370 appliances.

Figure 5-12 Predefined Mail Flow Policies for Single Listener

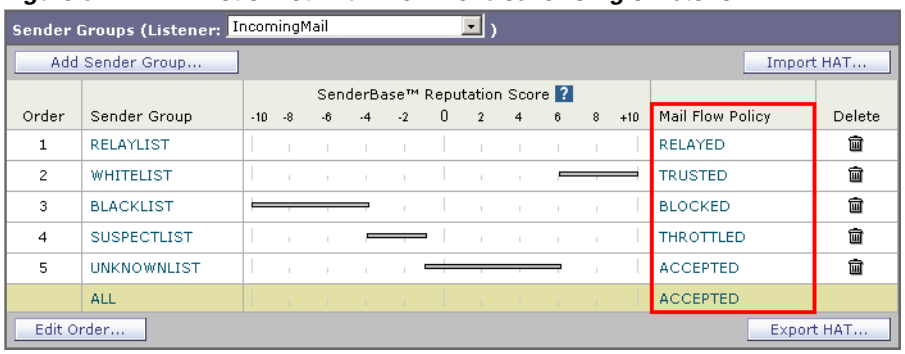

In this table, "Default" means that the default value as defined by the listener is used.

Table 5-11 Predefined Mail Flow Policies for Public Listeners

Policy Name	Primary Behavior (Access Rule)	Parameters	Value
$ACCEPTED (Used by All)	ACCEPT	Maximum messages / session: Maximum recipients / message: Maximum message size: Maximum concurrent connections: SMTP Banner Code: SMTP Banner Text: Override Hostname: Use TLS: Use McAfee virus scanning: Maximum recipients / hour: Maximum rcpt / hour Error Code: Max recipients / hour Text: Use SenderBase:	Default Default Default Default Default Default Default Default Default No default Default Default ON
Note: All parameters for the $ACCEPTED policy are user-defined in the CLI `systemsetup` and `listenerconfig` commands. Select "**y**" when prompted with the question: `Would you like to change the default host access policy?` to modify these values. To change these values using the GUI, follow the steps in Figure 5-7Viewing Default Mail Flow Policies, page 5-14.			
$BLOCKED	REJECT	Maximum messages / session: Maximum recipients / message: Maximum message size: Maximum concurrent connections: SMTP Banner Code: SMTP Banner Text: Override Hostname: Use TLS: Use McAfee virus scanning: Maximum recipients / hour: Maximum rcpt / hour Error Code: Max recipients / hour Text: Use SenderBase:	N/A N/A N/A N/A Default Default Default N/A N/A N/A N/A N/A N/A

Table 5-11 Predefined Mail Flow Policies for Public Listeners (Continued)

Policy Name	Primary Behavior (Access Rule)	Parameters	Value
$THROTTLED	ACCEPT	Maximum messages / session: Maximum recipients / message: Maximum message size: Maximum concurrent connections: SMTP Banner Code: SMTP Banner Text: Override Hostname Use TLS: Use McAfee virus scanning: Maximum recipients / hour: Maximum rcpt / hour Error Code: Max recipients / hour Text: Use SenderBase: Envelope Sender DNS Ver:	1 25 10MB 1 Default Default Default Default Default* 20 Default Default ON ON
$TRUSTED	ACCEPT	Maximum messages / session: Maximum recipients / message: Maximum message size: Maximum concurrent connections: SMTP Banner Code: SMTP Banner Text: Override Hostname: Use TLS: Use McAfee virus scanning: Maximum recipients / hour: Maximum rcpt / hour Error Code: Max recipients / hour Text: Use SenderBase:	5,000 5,000 100 MB 600 Default Default Default Default OFF* -1(Disable) N/A N/A OFF

* If enabled.

$ACCEPTED is a named policy, which is the same as the public listener's default HAT settings. You can assign the $ACCEPTED policy to any sender group you create. (See Adding a New Sender Group, page 5-31 and Connections, page 5-9. See also Working with the HAT, page 5-38).

The final ALL entry in a HAT for a public listener also uses the $ACCEPTED policy as the primary behavior.

Each public listener, by default, has the sender groups and corresponding mail flow policies shown in Table 5-12 defined by default.

Table 5-12 Predefined Sender Groups and Mail Flow Policies for Public Listeners

This Sender Group:	Uses this Mail Flow Policy:
WHITELIST	$TRUSTED
BLACKLIST	$BLOCKED
SUSPECTLIST	$THROTTLED
UNKNOWNLIST	$ACCEPTED

These four basic sender groups and mail flow policies enable a framework for you to begin classifying the email flowing into your gateway on a public listener. In "Using Email Security Monitor" in the *Cisco IronPort AsyncOS for Email Daily Management Guide*, you will be able to see the real-time flow of email

into your gateway and be able to make changes to a listener's HAT in real-time. (You can add IP addresses, domains, or organizations to an existing sender group, edit the existing or pre-defined policies, or create new mail flow policies.)

WHITELIST

Add senders you trust to the Whitelist sender group. The $TRUSTED mail flow policy is configured so that email from senders you trust has no rate limiting enabled, and the content from those senders is not scanned by the Anti-Spam or Anti-Virus software.

BLACKLIST

Senders in the Blacklist sender group are rejected (by the parameters set in the $BLOCKED mail flow policy). Adding senders to this group rejects connections from those hosts by returning a 5XX SMTP response in the SMTP HELO command.

SUSPECTLIST

The Suspectlist sender group contains a mail flow policy that throttles, or slows, the rate of incoming mail. If senders are suspicious, you can add them to the Suspectlist sender group, where the mail flow policy dictates that:

- Rate limiting limits the maximum number of messages per session, the maximum number of recipients per message, the maximum message size, and the maximum number of concurrent connections you are willing to accept from a remote host.
- The maximum recipients per hour from the remote host is set to 20 recipients per hour. Note that this setting is the maximum throttling available. You can increase the number of recipients to receive per hour if this parameter is too aggressive.
- The content of messages will be scanned by the anti-spam scanning engine and the anti-virus scanning engine (if you have these feature enabled for the system).
- The Cisco IronPort SenderBase Reputation Service will be queried for more information about the sender.

UNKNOWNLIST

The Unknownlist sender group may be useful if you are undecided about the mail flow policy you should use for a given sender. The mail flow policy for this group dictates that mail is accepted for senders in this group, but the Cisco IronPort Anti-Spam software (if enabled for the system), the anti-virus scanning engine, and the Cisco IronPort SenderBase Reputation Service should all be used to gain more information about the sender and the message content. Rate limits for senders in this group are also enabled with default values. For more information on virus scanning engines, see Anti-Virus Scanning, page 8-1. For more information on the SenderBase Reputation Service, see Reputation Filtering, page 7-1.

Predefined Mail Flow Policies for Private Listeners

When combined with an access rule (RELAY or REJECT), the parameters listed in Table 5-3 are predefined as the following two mail flow policies for each *private* listener you create:

- $RELAYED

- $BLOCKED

These policies are summarized in Table 5-12.

Figure 5-13 Predefined Mail Flow Policies for a Private Listener

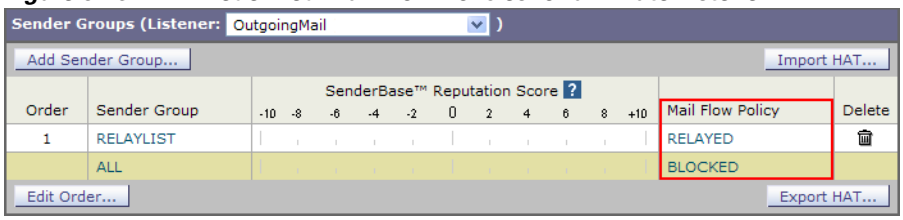

Table 5-13 Predefined Mail Flow Policies for Private Listeners

Policy Name	Primary Behavior (Access Rule)	Parameters	Value
$RELAYED	RELAY	Maximum messages / session: Maximum recipients / message: Maximum message size: Maximum concurrent connections: SMTP Banner Code: SMTP Banner Text: Override Hostname: Use TLS: Use Sophos virus scanning: Maximum recipients / hour: Maximum rcpt / hour Error Code: Max recipients / hour Text: Use SenderBase:	Default Default Default Default Default Default Default Default Off (if enabled) -1 (Disabled) Not applicable Not applicable Default
$BLOCKED (Used by All)	REJECT	Maximum messages / session: Maximum recipients / message: Maximum message size: Maximum concurrent connections: SMTP Banner Code: SMTP Banner Text: Override Hostname: Use TLS: Use Sophos virus scanning: Maximum recipients / hour: Maximum rcpt / hour Error Code: Max recipients / hour Text: Use SenderBase:	Not applicable Not applicable Not applicable Not applicable Default Default Default Not applicable Not applicable Not applicable Not applicable Not applicable Not applicable

$BLOCKED is a named policy, which is the same as the private listener's default HAT settings. The final ALL entry in a HAT for a private listener also uses the $BLOCKED policy as the default behavior.

Each private listener, by default, has the following predefined sender group and corresponding mail flow policy shown in Table 5-14 defined by default:

Table 5-14 Predefined Sender Groups and Mail Flow Policies for Private Listener

This Sender Group:	Uses this Mail Flow Policy:
RELAYLIST	$RELAYED
ALL	$BLOCKED

This basic sender group and mail flow policy enables a framework for you to begin classifying the email flowing out of your gateway on a private listener.

RELAYLIST

Add senders you know should be allowed to relay to the Relaylist sender group. The $RELAYED mail flow policy is configured so that email from senders you are allowing to relay has no rate limiting, and the content from those senders is not scanned by the anti-spam scanning engine or anti-virus software.

Note The systems you allowed to relay email through the Cisco IronPort appliance when you created an outbound (private) listener in the GUI System Setup wizard (or CLI `systemsetup` command) are automatically added to the RELAYLIST sender group. See Step 3: Network, page 3-17.

Note By default, C10/100 customers are prompted to create only one public listener during the `systemsetup` command. Public listeners created on Cisco IronPort C10/100 appliances also include a $RELAYED mail flow policy that is used to relay mail for internal systems.

Managing Sender Groups and Mail Flow Policies via the GUI

The Mail Policies > HAT Overview and Mail Flow Policy pages allow you to configure a HAT settings for a listener. From these pages, you can:

- See the mapping of sender groups to mail flow policies.
- Create, edit, or delete sender groups.
- Create, edit, or delete mail flow policies.
- Re-order HAT entries for a listener.

Click the Mail Policies > HAT Overview link. See Figure 5-14. Choose the listener you want to configure from the Listener: drop-down list.

Figure 5-14 Host Access Table Overview Page

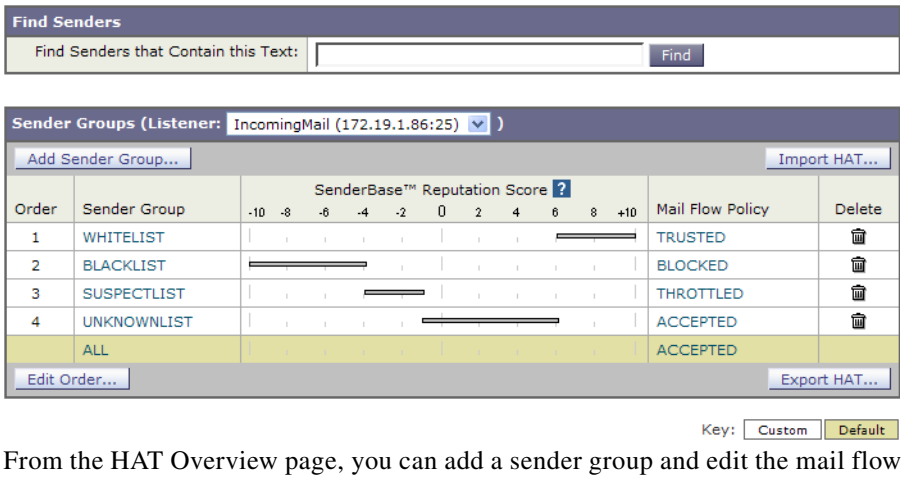

From the HAT Overview page, you can add a sender group and edit the mail flow policies for a listener.

Adding a New Sender Group

Step 1 From the HAT Overview page, click **Add Sender Group**.

Figure 5-15 Add Sender Group Page

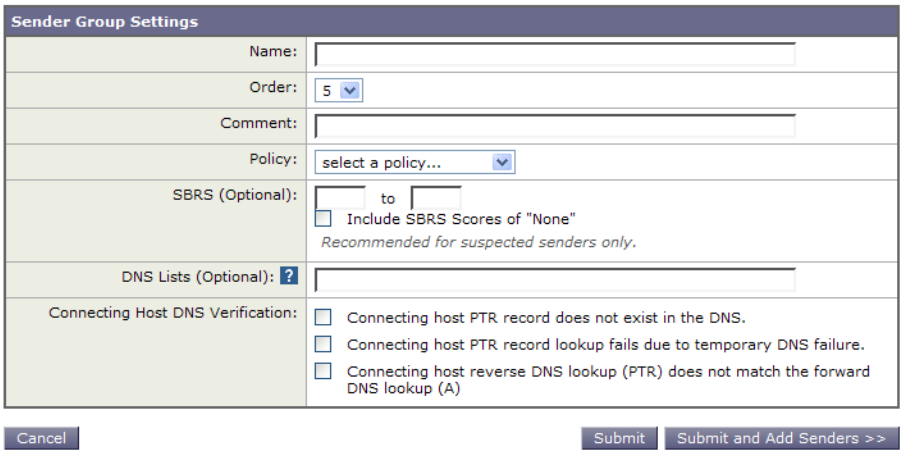

Step 2 Type the name of the sender group, select the order in which it will be placed in the list of sender groups, and a comment (optional) in the fields provided.

Step 3 If you do not know the mail flow policy you would like to apply to this group (or if no mail flow policies exist yet), then use the default "CONTINUE (no policy)" mail flow policy. Otherwise, choose a mail flow policy from the drop-down list.

Step 4 Select a SBRS range and DNS list (optional). You can also mark the checkbox to include senders for which SBRS has no information. This is referred to as "none" and generally denotes a suspect.

Step 5 Configure any host DNS verification settings (see Implementing Sender Verification — Example Settings, page 5-43).

Step 6 Click **Submit** to save the sender group and return to the Host Access Table page, or... click **Submit and Add Senders** to create the group and begin adding senders to it.

Step 7 Commit your changes.

Note If you attempt to enter duplicate entries (identical domain or IP addresses) in a single sender group, the duplicates are discarded.

Editing a Sender Group

Step 1 From the HAT Overview page, click the name of an existing sender group. The selected sender group is displayed:

Figure 5-16 Sender Group Detail Page

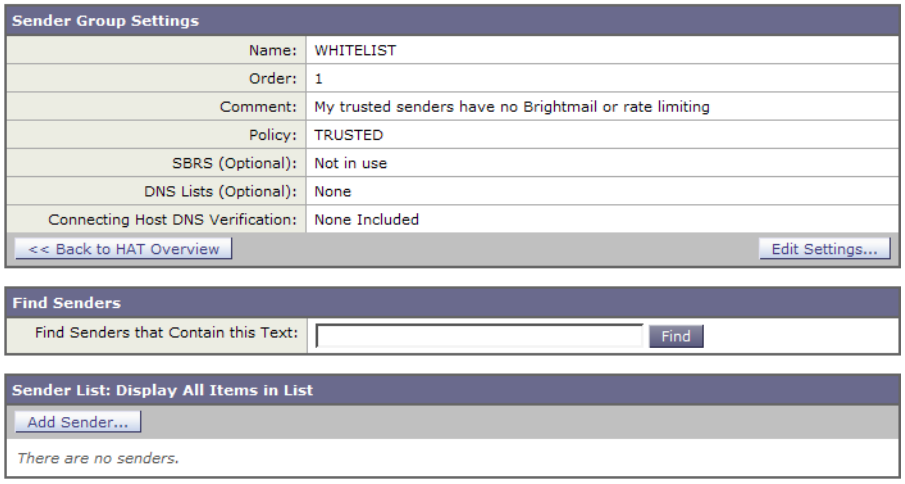

Step 2 Click **Edit Settings** The Edit Sender Group page is displayed.

Figure 5-17 Edit Sender Group Page

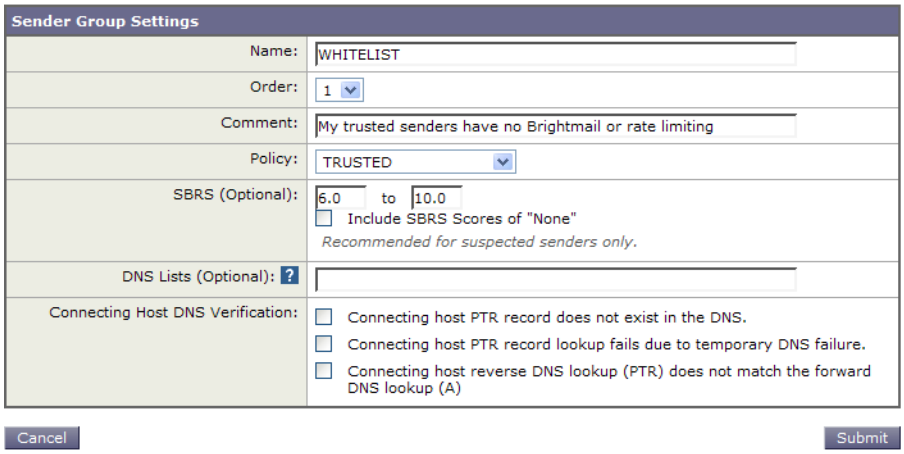

Step 3 Make changes to the sender group and click **Submit**.

Step 4 Commit your changes.

Deleting a Sender Group

Step 1 From the HAT Overview page, click the trash can icon in the Delete column for the sender group to delete. You are prompted to confirm the deletion.

Step 2 Click **Yes** to delete the sender group, or click **No** to cancel.

Step 3 Commit your changes.

Adding a New Mail Flow Policy

Step 1 Click the Mail Policies > Mail Flow Policies link. The Mail Flow Policies page is displayed.

Step 2 Click **Add Policy**. The Mail Flow Policies Add Policy page is displayed.

Step 3 Enter the information for the Mail Flow Policy.

Step 4 Configure envelope sender DNS verification settings (see Implementing Sender Verification — Example Settings, page 5-43).

Step 5 Submit and commit your changes.

Note Defaults for the policy are "greyed out" while the "Use Default" radio button is selected. To overwrite the default values, enable the feature or setting by selecting the "On" radio button and making changes to the now accessible values.

Note The Custom SMTP Banner Text and Max. Recipients Per Hour text string fields support the HAT variables discussed in HAT Variable Syntax, page 5-13.

Note Some parameters depend on certain pre-configurations. (For example, the Directory Harvest Attack prevention setting requires that you have configured an LDAP Acceptance Query.)

Editing a Mail Flow Policy

Step 1 From the Mail Flow Policy overview page, click the name of a policy. The Mail Flow Policy Edit Policy page is displayed.

Step 2 Make changes to the policy.

Step 3 Submit and commit your changes.

Deleting a Mail Flow Policy

Step 1 Click the trash can icon in the delete column for the mail flow policy to delete. You are prompted to confirm the deletion.

Step 2 Click **Yes** to delete the mail flow policy, or click **No** to cancel.

Step 3 Commit your changes.

Adding a Sender to a Sender Group

Step 1 From a domain, IP, or network owner profile page, click the Add to Sender Group link.

Figure 5-18 Add to Sender Group Link on a Profile Page

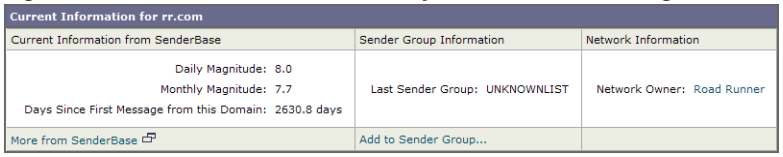

The Add to Sender Group page is displayed. See Figure 5-19.

Figure 5-19 Add to Sender Group Page

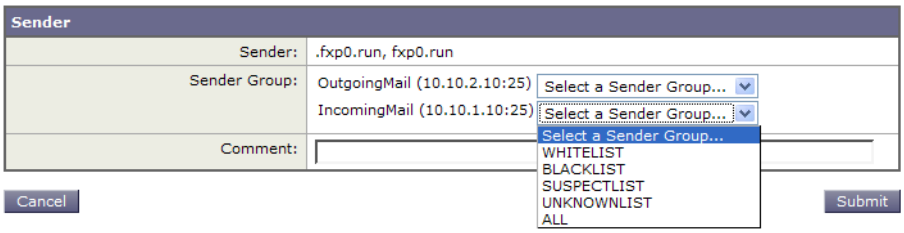

Step 2 Choose the sender group from the list defined for each listener.

Step 3 Click **Submit** to add the domain to the selected sender groups, or click **Cancel**.

Step 4 Commit your changes.

Note When you add a domain to a sender group, two actual domains are listed in the GUI. For example, if you were adding the domain `example.net`, on the Add to Sender Group page, both `example.net` and `.example.net` are added. The second entry ensures that any host in the subdomain of `example.net` will be added to the sender group. For more information, see Sender Group Syntax, page 5-21.

Note If one or more of the senders you are adding to a sender group is a duplicate of a sender that is already present in that sender group, the duplicate senders will not be added and you will see a confirmation message.

Success — Added sender(s) to sender group(s). Some duplicates existed and were not added.

Step 5 Click **Save** to add the sender and return to the Incoming Mail Overview page.

Adding a Sender to a New Sender Group

Step 1 When creating a new Sender Group, click **Submit and Add Senders**. The Add Sender page is displayed.

Step 2 Enter a sender using an IPv4 address, an IPv6 address, or a hostname. A sender can include a range of IP addresses and partial hostnames.

Step 3 Enter an optional comment for the sender.

Step 4 Click **Submit** to add the domain to the sender group, or click **Cancel**.

Step 5 Commit your changes.

Creating a Sender Group with SenderBase Reputation Scores

Step 1 Click **Add Sender Group** from the HAT Overview page.

Step 2 On the Add Sender Group page, type the name of the sender group and an optional comment.

Step 3 Choose a mail flow policy from the list.

Step 4 In the Senders section, choose SBRS from the drop-down list and click **Add Sender**.

The page refreshes.

Step 5 Type the range in the SBRS from: and to: fields, and an optional comment.

In Figure 5-20, senders with a SenderBase Reputation Score less than -7.5 are blocked using the BLOCKED mail flow policy.

Figure 5-20 Creating a Sender Group with SenderBase Reputation Scores (1)

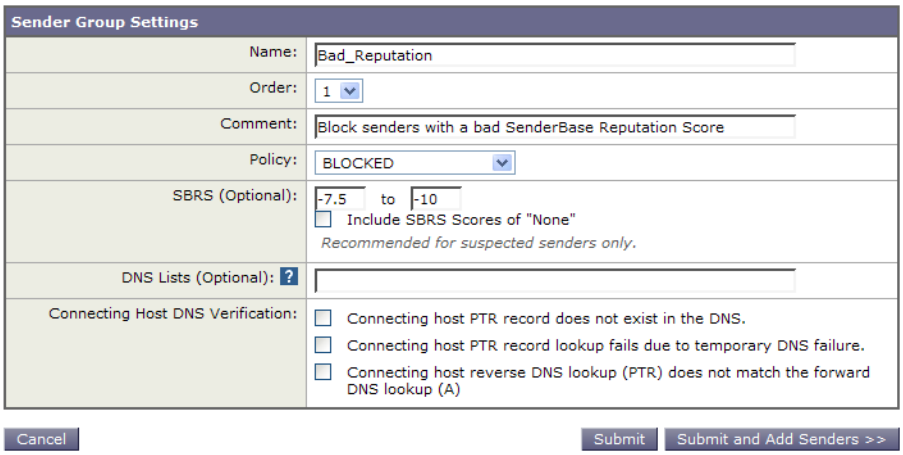

In Figure 5-21, senders with a SenderBase Reputation Score greater than 8.0 bypass the anti-spam scanning for the listener:

Figure 5-21 Creating a Sender Group with SenderBase Reputation Scores (2)

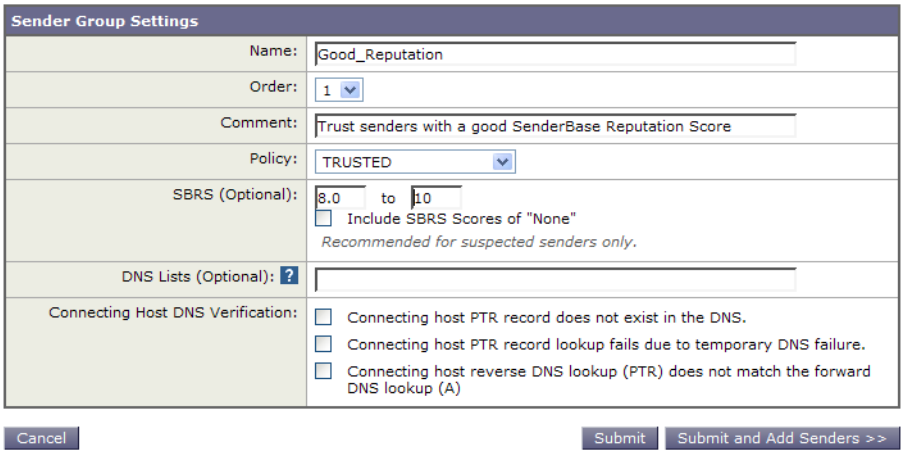

Note You can also modify the default policies of the TRUSTED and BLOCKED to include senders based on SenderBase Reputation Scores using these same parameters. See Implementing SenderBase Reputation Filters, page 7-4 for more information.

Step 6 Click **Submit** to create the sender group based on SenderBase Reputation Scores.

Step 7 Commit your changes.

Figure 5-22 Host Access Table Using SenderBase Reputation Scores

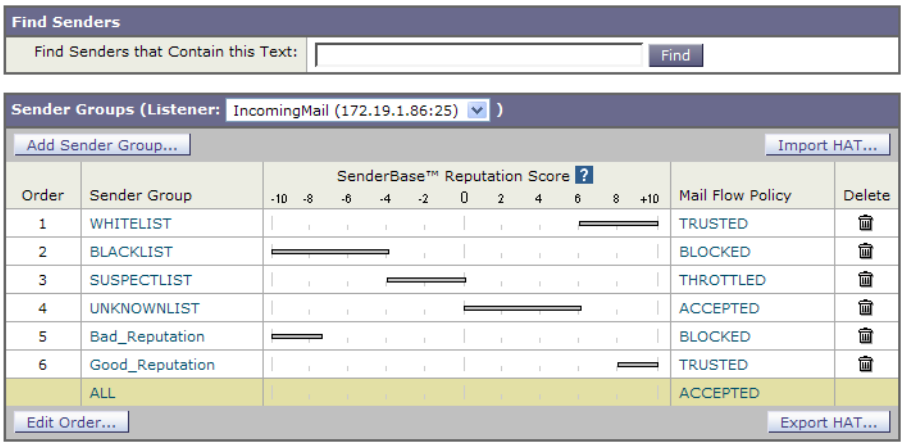

Reordering the HAT

The order of entries in a HAT is important. Remember that the HAT is read from top to bottom for each host that attempts to connect to the listener. If a rule matches a connecting host, the action is taken for that connection immediately.

For example, if you specify a CIDR block in Sender Group A (using policy 1) and create Sender Group B for an IP address within that CIDR block, the policy in Sender Group B will never be applied.

Step 1 From the HAT Overview page, click **Edit Order** The Edit Sender Group Order page is displayed.

Step 2 Type the new order for existing rows of the HAT.

Step 3 Submit and commit your changes.

The HAT Overview page refreshes with the new order displayed.

In the following example shown in Figure 5-23, the order is being changed so that trusted senders are processed first, blocked senders are processed next, and unknown or suspected senders are processed last.

Figure 5-23 Changing the Order of Entries in the HAT

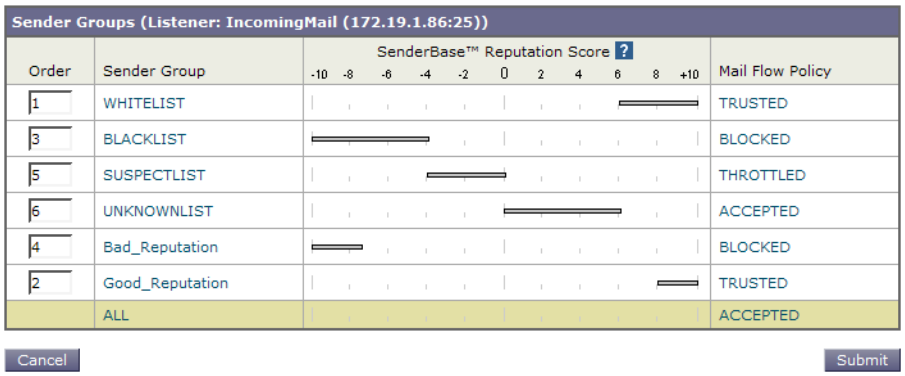

Searching for Senders

You can find senders by entering text in the Find Senders field at the top of the HAT Overview page. Enter the text to search with and click Find.

Modifying the HAT for a Listener via the GUI

Log in to the Graphical User Interface (GUI) and click the Mail Policies tab. (For information about how to access the GUI, see Accessing the GUI, page 2-2.) Click the HAT Overview link in the left menu. The Host Access Table Overview page is displayed:

Figure 5-24 The Host Access Table Overview Page

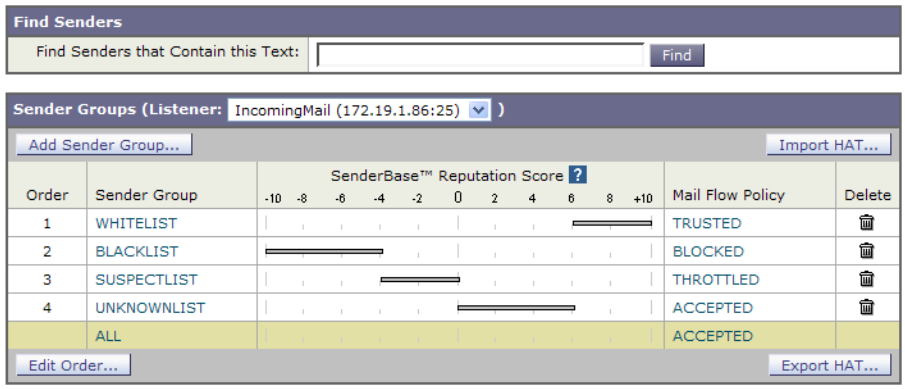

The Host Access Table Overview shows a listing of the sender groups in the HAT, including the order, SenderBase Reputation Score range, and associated Mail Flow Policy.

From the Host Access Table Overview, you can:

- Add sender groups to the HAT
- Delete sender groups from the HAT
- Modify existing sender groups
- Change the order of the entries

- Import a HAT (overwrites existing entries) from a file (importing and exporting the HAT is described below, see Working with the HAT, page 5-38)
- Export the HAT to a file
- Search for senders

Once you are editing a sender group, you can:

- Add senders to (and remove senders from) sender groups
- Edit settings for a sender group

For more information about working with Sender Groups, see Managing Sender Groups and Mail Flow Policies via the GUI, page 5-30.

Working with the HAT

Exporting the HAT

Step 1 Click **Export HAT** The Export Host Access Table page is displayed:

Figure 5-25 Exporting the HAT

Step 2 Enter a file name for the exported HAT. This is the name of the file that will be created in the configuration directory on the appliance.

Step 3 Submit and commit your changes.

Importing a HAT

When you import a HAT, all of the existing HAT entries are removed from the current HAT.

Step 1 Click **Import HAT** The Import Host Access Table page is displayed:

Figure 5-26 Exporting a HAT

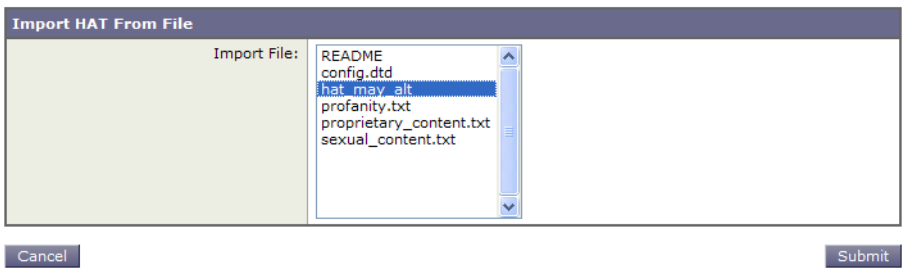

Step 2 Select a file from the list.

Note The file to import must be in the configuration directory on the appliance.

Step 3 Click **Submit**. You will see a warning message, asking you to confirm that you wish to remove all of the existing HAT entries.

Step 4 Click **Import**.

Step 5 Commit your changes.

You can place "comments" in the file. Lines that begin with a '#' character are considered comments and are ignored by AsyncOS. For example:

```
# File exported by the GUI at 20060530T215438

$BLOCKED

    REJECT {}

[ ... ]
```

Address Lists

Mail flow policies allow you to use of an address list for certain settings that apply to a group of envelope senders, such as rate limiting exemptions and mandatory TLS connections. An address list can consist of email addresses, domains, partial domains, and IP addresses. You can use the **Mail Policies > Address Lists** page in the GUI or the `addresslistconfig` command in the CLI to create an address list. The Address Lists page displays all address lists on the appliance, along with any mail flow policies that use an address list.

Creating an Address List

Step 1 Select **Mail Policies > Address Lists**.

Step 2 Click **Add Address List**.

The Add Address List page is displayed.

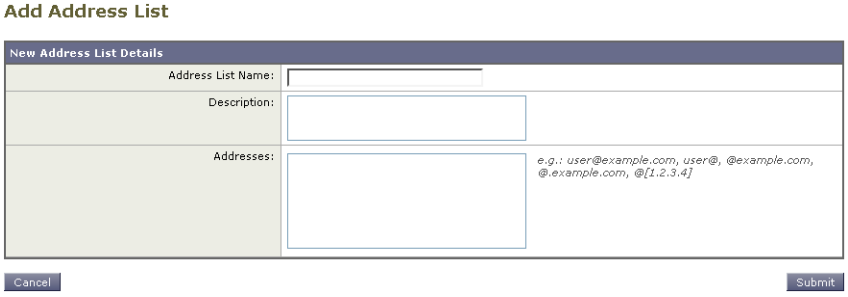

Step 3 Enter a name for the address list.

Step 4 Enter a description of the address list.

Step 5 Enter the addresses you want to include. You can use the following formats:

- Full email address: `user@example.com`

- Partial email address: `user@`
- All users in a domain: `@example.com`
- All users in a partial domain: `@.example.com`
- All users from a host with a certain IP address: `@[1.2.3.4]`

Note that domains and IP addresses must start with a `@` character.

Separate email addresses with a comma. If you separate the addresses using a new line, AsyncOS automatically converts your entries into a comma-separate list.

Step 6 Submit and commit your changes.

Editing an Address List

Step 1 Select **Mail Policies > Address Lists**.

Step 2 Click the name of the address list you want to edit.

Step 3 Modify the address list.

Step 4 Submit and commit your changes.

Deleting an Address List

To delete an address list, check the Delete check box next to the message action you want to delete. A confirmation message notifies you if any mail flow policies use the address list. Deleting a list removes it from the mail flow policies that use it. Commit your changes.

Sender Verification

Spam and unwanted mail is frequently sent by senders whose domains or IP addresses cannot be resolved by DNS. DNS verification means that you can get reliable information about senders and process mail accordingly. Sender verification prior to the SMTP conversation (connection filtering based on DNS lookups of the sender's IP address) also helps reduce the amount of junk email processed through the mail pipeline on the Cisco IronPort appliance.

Mail from unverified senders is not automatically discarded. Instead, AsyncOS provides sender verification settings that allow you to determine how the appliance handles mail from unverified senders: you can configure your Cisco IronPort appliance to automatically block all mail from unverified senders prior to the SMTP conversation or throttle unverified senders, for example.

The sender verification feature consists of two components: verification of the connecting host, which occurs prior to the SMTP conversation, and verification of the domain portion of the envelope sender, which occurs during the SMTP conversation.

Sender Verification: Host

Senders can be unverified for different reasons. For example, the DNS server could be "down" or not responding, or the domain may not exist. Host DNS verification settings for sender groups allow you to classify unverified senders prior to the SMTP conversation and include different types of unverified senders in your various sender groups.

The Cisco IronPort appliance attempts to verify the sending domain of the connecting host via DNS for incoming mail. This verification is performed prior to the SMTP conversation. The system acquires and verifies the validity of the remote host's IP address (that is, the domain) by performing a *double DNS lookup*. A double DNS lookup is defined as a reverse DNS (PTR) lookup on the IP address of the connecting host, followed by a forward DNS (A) lookup on the results of the PTR lookup. The appliance then checks that the results of the A lookup match the results of the PTR lookup. If the PTR or A lookups fail, or the results do not match, the system uses only the IP address to match entries in the HAT and the sender is considered as not verified.

Unverified senders are classified into three categories:

- Connecting host PTR record does not exist in the DNS.
- Connecting host PTR record lookup fails due to temporary DNS failure.
- Connecting host reverse DNS lookup (PTR) does not match the forward DNS lookup (A).

Using the sender group "Connecting Host DNS Verification" settings, you can specify a behavior for unverified senders (see Implementing Host Sender Verification for the SUSPECTLIST Sender Group, page 5-44).

You can enable host DNS verification in the sender group settings for any sender group; however, keep in mind that adding host DNS verification settings to a sender group means *including* unverified senders in that group. That means that spam and other unwanted mail will be included. Therefore, you should only enable these settings on sender groups that are used to reject or throttle senders. Enabling host DNS verification on the WHITELIST sender group, for example, would mean that mail from unverified senders would receive the same treatment as mail from your trusted senders in your WHITELIST (including bypassing anti-spam/anti-virus checking, rate limiting, etc., depending on how the mail flow policy is configured).

Sender Verification: Envelope Sender

With envelope sender verification, the domain portion of the envelope sender is DNS verified. (Does the envelope sender domain resolve? Is there an A or MX record in DNS for the envelope sender domain?) A domain does not resolve if an attempt to look it up in the DNS encounters a temporary error condition such as a timeout or DNS server failure. On the other hand, a domain does not exist if an attempt to look it up returns a definitive "domain does not exist" status. This verification takes place during the SMTP conversation whereas host DNS verification occurs before the conversation begins — it applies to the IP address of connecting SMTP server.

In more detail: AsyncOS performs an MX record query for the domain of the sender address. AsyncOS then performs an A record lookup based on the result of the MX record lookup. If the DNS server returns "NXDOMAIN" (there is no record for this domain), AsyncOS treats that domain as non-existent. This falls into the category of "Envelope Senders whose domain does not exist." NXDOMAIN can mean that the root name servers are not providing any authoritative name servers for this domain.

However, if the DNS server returns "SERVFAIL," it is categorized as "Envelope Senders whose domain does not resolve." SERVFAIL means that the domain does exist but DNS is having transient problems looking up the record.

A common technique for spammers or other illegitimate senders of mail is to forge the MAIL FROM information (in the envelope sender) so that mail from unverified senders that is accepted will be processed. This can lead to problems as bounce messages sent to the MAIL FROM address are undeliverable. Using envelope sender verification, you can configure your Cisco IronPort appliance to reject mail with malformed (but not blank) MAIL FROMs.

For each mail flow policy, you can:

- Enable envelope sender DNS verification.
- Offer custom SMTP code and response for malformed envelope sender. Malformed envelope senders are blocked if you have enabled envelope sender DNS verification.
- Offer custom response for envelope sender domains which do not resolve.
- Offer custom response for envelope sender domains which do not exist in DNS.

You can use the sender verification exception table to store a list of domains or addresses from which mail will be automatically allowed or rejected (see Sender Verification Exception Table, page 5-43). The sender verification exception table can be enabled independently of Envelope Sender verification. So, for example, you can still reject special addresses or domains specified in the exception table without enabling envelope sender verification. You can also always allow mail from internal or test domains, even if they would not otherwise be verified.

Though most spam is from unverifiable senders, there are reasons why you might want to accept mail from an unverified sender. For example, not all legitimate email can be verified through DNS lookups — a temporary DNS server problem can stop a sender from being verified.

When mail from unverified senders is attempted, the sender verification exception table and mail flow policy envelope sender DNS verification settings are used to classify envelope senders during the SMTP conversation. For example, you may accept and throttle mail from sending domains that are not verified because they do not exist in DNS. Once that mail is accepted, messages with malformed MAIL FROMs are rejected with a customizable SMTP code and response. This occurs during the SMTP conversation.

You can enable envelope sender DNS verification (including the domain exception table) in the mail flow policy settings for any mail flow policy via the GUI or the CLI (`listenerconfig -> edit -> hostaccess -> <policy>`).

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 (see the SMTP Address Parsing Options section in “Customizing Listeners” in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*), the default domain settings for that listener will no longer be used.

These features are mutually exclusive.

Custom SMTP Code and Response

You can specify the SMTP code and response message for messages with malformed envelope senders, for envelope senders which do not exist in DNS, and for envelope senders which do not resolve via DNS queries (DNS server might be down, etc.).

In the SMTP response, you can include a variable, `$EnvelopeSender`, which is expanded to the value of the envelope sender when the custom response is sent.

While typically a “Domain does not exist” result is permanent, it is possible for this to be a transient condition. To handle such cases, “conservative” users may wish to change the error code from the default 5XX to a 4XX code.

Sender Verification Exception Table

The sender verification exception table is a list of domains or email addresses that will either be automatically allowed or rejected during the SMTP conversation. You can also specify an optional SMTP code and reject response for rejected domains. There is only one sender verification exception table per Cisco IronPort appliance and it is enabled per mail flow policy.

The sender verification exception table can be used to list obviously fake but correctly formatted domains or email addresses from which you want to reject mail. For example, the correctly formatted MAIL FROM: `pres@whitehouse.gov` could be listed in the sender verification exception table and set to be automatically rejected. You can also list domains that you want to automatically allow, such as internal or test domains. This is similar to envelope recipient (SMTP RCPT TO command) processing which occurs in the Recipient Access Table (RAT).

The sender verification exception table is defined in the GUI via the Mail Policies > Exception Table page (or the CLI, via the `exceptionconfig` command) and then is enabled on a per-policy basis via the GUI (see Implementing Sender Verification for the ACCEPTED Mail Flow Policy, page 5-46) or the CLI (see the *Cisco IronPort AsyncOS CLI Reference Guide*.

Entries in the sender verification exception table have the following syntax:

Figure 5-27 Exception Table Listing

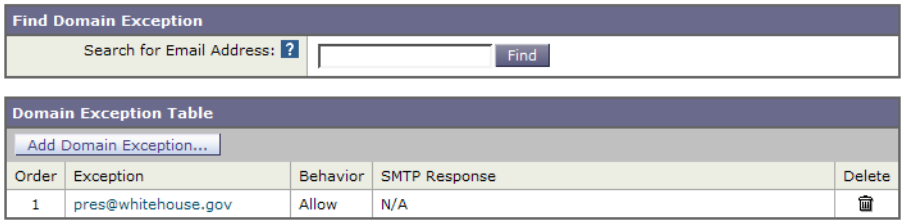

See Creating the Sender Verification Exception Table via the GUI, page 5-47 for more information about modifying the exception table.

Implementing Sender Verification — Example Settings

This section provides an example of a typical conservative implementation of host and envelope sender verification.

For this example, when implementing host sender verification, mail from connecting hosts for which reverse DNS lookup does not match is throttled via the existing SUSPECTLIST sender group and THROTTLED mail flow policy.

A new sender group (UNVERIFIED) and a new mail flow policy (THROTTLEMORE) are created. Mail from connecting hosts which are not verified will be throttled (using the UNVERIFIED sender group and the more aggressive THROTTLEMORE mail flow policy) prior to the SMTP conversation.

Envelope sender verification is enabled for the ACCEPTED mail flow policy.

Table 5-15 shows the suggested settings for implementing sender verification:

Table 5-15 Sender Verification: Suggested Settings

Sender Group	Policy	Include
UNVERIFIED	THROTTLEMORE	Prior to SMTP conversation: Connecting host PTR record does not exist in the DNS.
SUSPECTLIST	THROTTLED	Connecting host reverse DNS lookup (PTR) does not match the forward DNS lookup (A).
	ACCEPTED	Envelope Sender Verification during SMTP conversation: - Malformed MAIL FROM: - Envelope sender does not exist in DNS. - Envelope sender DNS does not resolve.

Implementing Host Sender Verification for the SUSPECTLIST Sender Group

Step 1 Select **Mail Policies > HAT Overview**.

Step 2 Click **SUSPECTLIST** in the list of sender groups.

Figure 5-28 HAT Overview Page

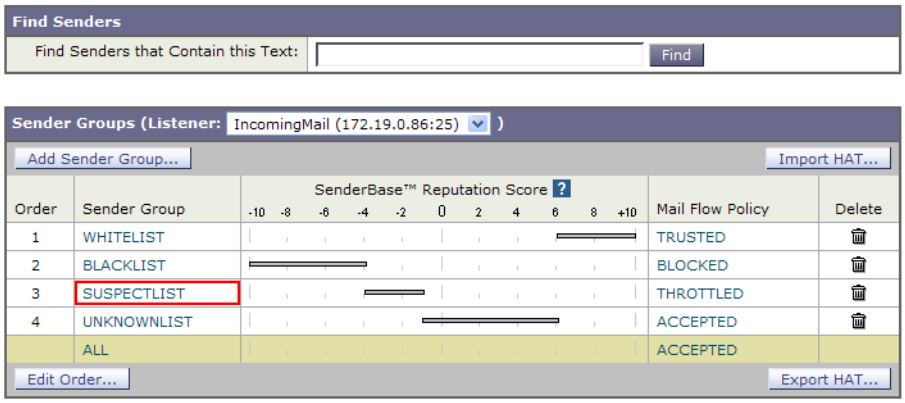

Step 3 The Sender Group: SUSPECTLIST page is displayed:

Figure 5-29 Sender Group: SUSPECTLIST

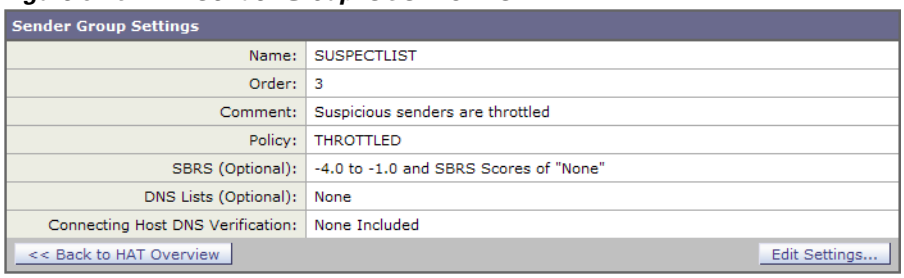

Step 4 Click **Edit Settings** The Edit Settings dialog is displayed:

Figure 5-30 Sender Group: SUSPECTLIST: Edit Settings

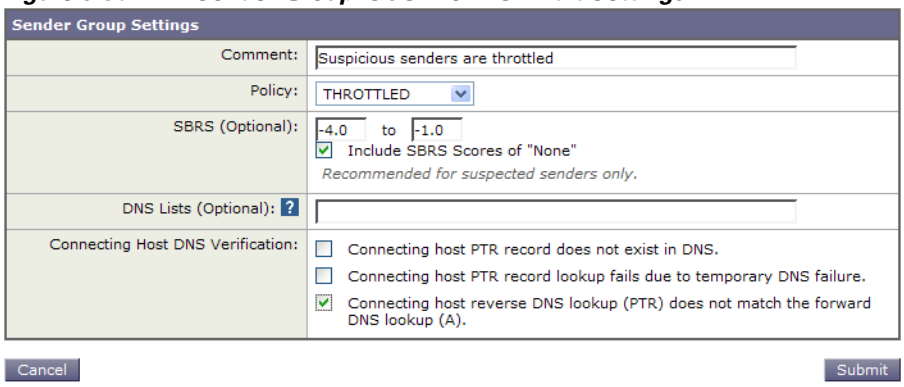

Step 5 Select the THROTTLED policy from the list.

Step 6 Check the "Connecting host reverse DNS lookup (PTR) does not match the forward DNS lookup (A)" checkbox under Connecting Host DNS Verification.

Step 7 Submit and commit your changes.

Now, senders for which reverse DNS lookups fail will match the SUSPECTLIST sender group and will receive the default action from the THROTTLED mail flow policy.

Note You can also configure host DNS verification via the CLI. See Enabling Host DNS Verification via the CLI, page 5-50 for more information.

Implementing Sender Verification

First, create a new mail flow policy (for this example, it is named THROTTLEMORE) and configure it with more stringent throttling settings.

Step 1 On the Mail Flow Policies page, click **Add Policy**

Step 2 Enter a name for the mail flow policy, and select Accept as the Connection Behavior.

Step 3 Configure the policy to throttle mail.

Step 4 Submit and commit your changes.

Next, create a new sender group (for this example, it is named UNVERIFIED) and configure it to use the THROTTLEMORE policy:

Step 1 On the HAT Overview page, click **Add Sender Group**

Figure 5-31 Add Sender Group: THROTTLEMORE

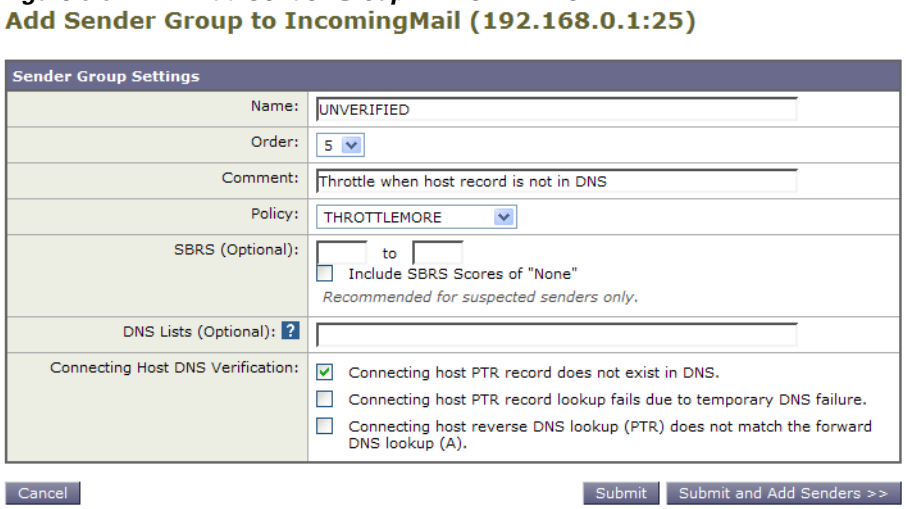

Step 2 Select the THROTTLEMORE policy from the list.

Step 3 Check the "Connecting host PTR record does not exist in DNS" checkbox under Connecting Host DNS Verification.

Step 4 Submit and commit your changes. The HAT Overview page now looks like this:

Figure 5-32 HAT Overview

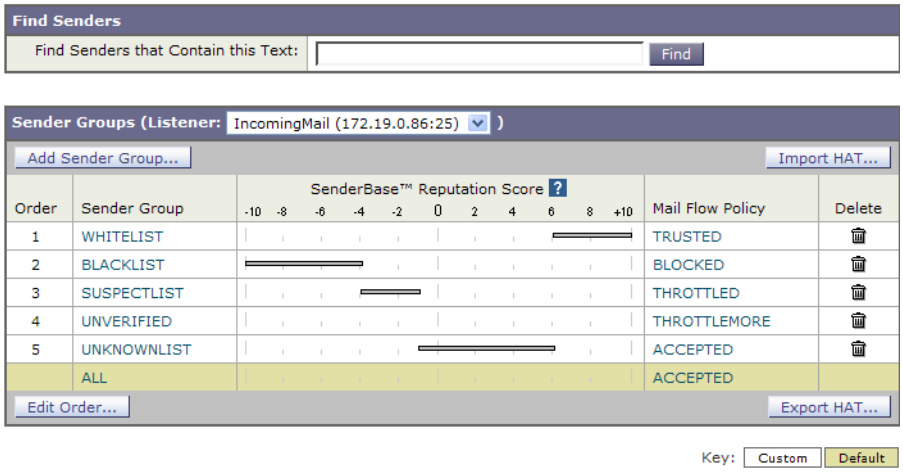

For the next step, configure the ACCEPTED mail flow policy to handle unverified senders.

Implementing Sender Verification for the ACCEPTED Mail Flow Policy

Step 1 Select **Mail Policies > Mail Flow Policies**.

Step 2 On the Mail Flow Policies page, click on the ACCEPTED mail flow policy.

Step 3 Scroll to the bottom of the mail flow policy:

Figure 5-33 ACCEPTED Mail Flow Policy Envelope Sender DNS Verification Settings

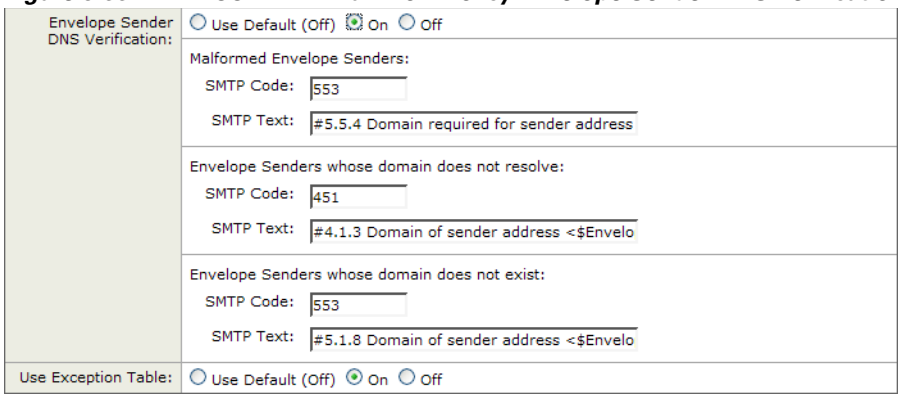

Step 4 Select On to enable envelope sender DNS verification for this mail flow policy.

Step 5 You may also define custom SMTP code and responses.

Step 6 Enable the domain exception table by selecting On for "Use Domain Exception Table."

Step 7 Submit and commit your changes.

And for the last step, create the sender verification exception table to list exceptions to the sender verification settings.

Creating the Sender Verification Exception Table via the GUI

Step 1 Select **Mail Policies > Exception Table**.

Note The exception table applies globally to all mail flow policies with "Use Exception Table" enabled.

Step 2 Click **Add Domain Exception** on the Mail Policies > Exception Table page. The Add Domain Exception page is displayed:

Figure 5-34 Adding Addresses to the Exception Table

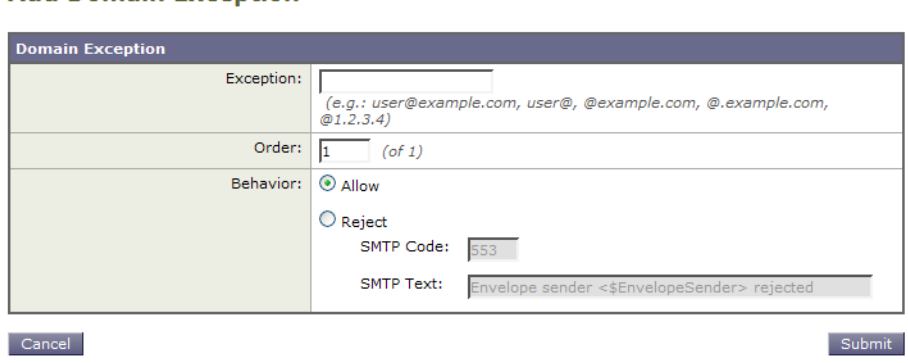

Step 3 Enter an email address. You can enter a specific address (pres@whitehouse.gov), a name (user@), a domain (@example.com or @.example.com), or an address with a bracketed IP address (user@[192.168.23.1]).

Step 4 Specify whether to allow or reject messages from the address. When rejecting mail, you can also specify an SMTP code and custom response.

Step 5 Submit and commit your changes.

Searching for Addresses within the Sender Verification Exception Table

Step 1 Enter the email address in the Find Domain Exception section of the Exception Table page and click **Find**.

Figure 5-35 Searching for Matching Entries in the Exception Table

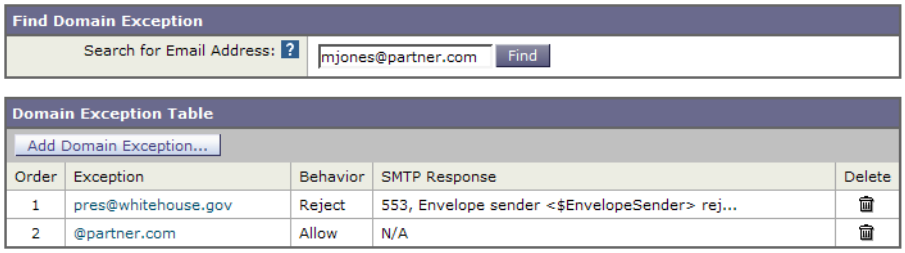

Step 2 If the address matches any of the entries in the table, the first matching entry is displayed:

Figure 5-36 Listing Matching Entries in the Exception Table

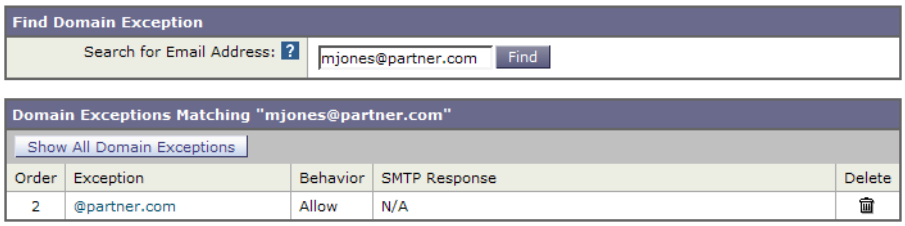

Testing Sender Verification Settings

Now that you have configured sender verification settings, you can verify the behavior of your Cisco IronPort appliance.

Note that testing DNS-related settings is beyond the scope of this document.

Testing the Envelope Sender Verification Settings

While it may be difficult to test the various DNS-related settings for your THROTTLED policy, you can test the malformed MAIL FROM setting.

Step 1 Open a Telnet session to your Cisco IronPort appliance.

Step 2 Use SMTP commands to send a test message with a malformed MAIL FROM (something like “admin” without a domain).

Note If you have configured your Cisco IronPort appliance to use a default domain or to specifically allow partial domains when sending or receiving email or if you have enabled address parsing (see "Customizing Listeners" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*) you may not be able to create, send, and receive an email with a missing or malformed domain.

Step 3 Verify that the message is rejected.

```
# telnet IP_address_of_IronPort_Appliance port

220 hostname ESMTP

helo example.com

250 hostname

mail from: admin

553 #5.5.4 Domain required for sender address
```

Note that the SMTP code and response is the one you configured for the envelope sender verification settings for the THROTTLED mail flow policy.

Testing the Sender Verification Exception Table

To confirm that mail from the email address listed in the sender verification exception table is not subject to envelope sender verification:

Step 1 Add the following address to the exception table with an "Allow" behavior: admin@zzzaaazzz.com

Step 2 Commit your changes.

Step 3 Open a Telnet session to your Cisco IronPort appliance.

Step 4 Use SMTP commands to send a test message from the email address you entered in the sender verification exception table (admin@zzzaaazzz.com).

Step 5 Verify that the message is accepted.

```
# telnet IP_address_of_IronPort_Appliance port

220 hostname ESMTP

helo example.com

250 hostname

mail from: admin@zzzaaazzz.com

250 sender <admin@zzzaaazzz.com> ok
```

If you remove that email address from the sender verification exception table, mail from that sender will be rejected because the domain portion of the envelope sender is not DNS verified.

Sender Verification and Logging

The following log entries provide an example of Sender Verification verdicts.

Envelope Sender Verification

Malformed Envelope Senders:

```
Thu Aug 10 10:14:10 2006 Info: ICID 3248 Address: <user> sender rejected, envelope
sender domain missing
```

Domain does not exist (NXDOMAIN):

```
Wed Aug  9 15:39:47 2006 Info: ICID 1424 Address: <user@domain.com> sender rejected,
envelope sender domain does not exist
```

Domain does not resolve (SERVFAIL):

```
Wed Aug  9 15:44:27 2006 Info: ICID 1425 Address: <user@domain.com> sender rejected,
envelope sender domain could not be resolved
```

Enabling Host DNS Verification via the CLI

To enable host DNS verification in the CLI, use the `listenerconfig->edit->hostaccess` command (see the *Cisco IronPort AsyncOS CLI Reference Guide* for more information).

Table 5-16 shows the types of unverified senders and the corresponding CLI setting:

Table 5-16 Sender Group Settings and Corresponding CLI Values

Connecting Host DNS Verification	Equivalent CLI Setting
Connecting host PTR record does not exist in the DNS.	`nx.domain`
Connecting host PTR record lookup fails due to temporary DNS failure.	`serv.fail`
Connecting host reverse DNS lookup (PTR) does not match the forward DNS lookup (A)	`not.double.verified`

Accepting Email for Local Domains or Specific Users on Public Listeners (RAT)

When you create a public listener, you define all local domains that the appliance will accept messages for using the Recipient Access Table (RAT). Many enterprise gateways are configured to receive messages for several local domains. For example, suppose your company changed its name. You would need to receive email messages for recipients addressed to `currentcompanyname.com` and `oldcompanyname.com`. In this case, both local domains would be included in the RAT for your public

listener. (Note: the Domain Map feature can map messages from one domain to another. See the Domain Map feature section of the "Configuring Routing and Domain Features" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.)

Note If you have completed System Setup Wizard or the `systemsetup` command and issued the `commit` command, one public listener should already be configured on your appliance. (Refer to the settings you entered for: Step 3: Network, page 3-17.) The default local domains or specific addresses to accept mail that you entered at that time were the first entries in the RAT for that public listener.

Recipient Access Table (RAT)

The Recipient Access Table defines which recipients will be accepted by the public listener. The table specifies the address (which may be a partial address, username, domain, or hostname) and whether to accept or reject it. You can optionally include the SMTP response to the RCPT TO command for that recipient or bypass throttling control for specific entries.

RAT entries are defined by this basic syntax:

Table 5-17 Basic RAT Syntax

Recipient Definition	Rule	(Optional) Custom SMTP Response

Rules

The RAT has two basic actions that it performs on recipients as they communicate in the SMTP conversation:

`ACCEPT`	The recipient is accepted.
`REJECT`	The recipient is rejected.

Defining Recipients

The RAT allows you to define a recipient or group of recipients. Recipients can be defined by full email address, domain, partial domain, username, or IP address:

[IPv4 address]	Specific Internet Protocol version 4 (IPv4) address of the host. Note that the IP address must be between the "`[]`" characters.
[IPv6 address]	Specific Internet Protocol version 6 (IPv6) address of the host. Note that the IP address must be between the "`[]`" characters.
`division.example.com`	Fully-qualified domain name.
`.partialhost`	Everything within the "partialhost" domain.
`user@domain`	Complete email address.

user@	Anything with the given username.
user@[*IP_address*]	Username at a specific IPv4 or IPv6 address. Note that the IP address must be between the "`[]`" characters. Note that "`user@`*IP_address*" (without the bracket characters) is not a valid address. The system will append the brackets when it receives the message to create a valid address, which could affect whether a recipient is matched in the RAT.

Note When you add a domain to the Recipient Access Table in step 4 of the System Setup Wizard in the GUI (see Step 3: Network, page 3-17), you might want to consider adding a second entry to specify subdomains. For example, if you type the domain `example.net`, you might also want to enter `.example.net`. The second entry ensures that mail destined for any subdomain of `example.net` will match in the Recipient Access Table. Note that *only* specifying `.example.com` in the RAT will accept for all subdomains of `.example.com` but *will not* accept mail for complete email address recipients *without* a subdomain (for example joe@example.com).

Bypassing Throttling for Special Recipients

For recipient entries, you can specify that the recipient bypasses throttling control mechanisms enabled on the listener.

This feature is useful if there are certain recipients for whom you do not want to limit messages. For example, many users will want to receive email for the address "`postmaster@domain`" on a listener, even if the sending domain is being throttled based on the receiving control defined in mail flow policies. Specifying this recipient to bypass receiving control in a listener's RAT allows the listener to receive unlimited messages for the recipient "`postmaster@domain`" while retaining mail flow policies for other recipients in the same domain. Recipients will avoid being counted against the recipients-per-hour counter maintained by the system if the sending domain is being limited.

To specify certain recipients to bypass receiving control via the GUI, select "Yes" for Bypass Receiving Control when adding or editing a RAT entry:

Figure 5-37 Bypassing Receiving Control

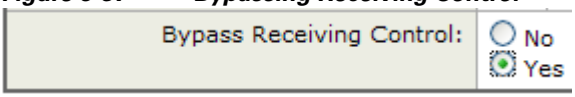

To specify certain recipients to bypass receiving control via the CLI, answer yes to the following question when you enter recipients using the `listenerconfig -> edit -> rcptaccess` command:

```
Would you like to bypass receiving control for this entry?  [N]> y
```

Bypassing LDAP Accept for Special Recipients

If you configure LDAP acceptance queries, you may wish to bypass the acceptance query for certain recipients. This feature can be useful if there are recipients for whom you receive email which you do not want to be delayed or queued during LDAP queries, such as `customercare@example.com`.

If you configure the recipient address to be rewritten in the work queue prior to the LDAP acceptance query, (such as aliasing or using a domain map), the rewritten address will not bypass LDAP acceptance queries. For example you use an alias table to map `customercare@example.com` to `bob@example.com` and `sue@example.com`. If you configure bypassing LDAP acceptance for `customercare@example.com`, an LDAP acceptance query is still run for `bob@example.com` and `sue@example.com` after the aliasing takes place.

To configure bypassing LDAP acceptance via the GUI, select **Bypass LDAP Accept Queries for this Recipient** when you add or edit the RAT entry.

To configure bypassing LDAP acceptance queries via the CLI, answer yes to the following question when you enter recipients using the `listenerconfig -> edit -> rcptaccess` command:

```
Would you like to bypass LDAP ACCEPT for this entry? [Y]> y
```

When you configure a RAT entry to bypass LDAP acceptance, be aware that the order of RAT entries affects how recipient addresses are matched. The RAT matches the recipient address with the first RAT entry that qualifies. For example, you have the following RAT entries: postmaster@ironport.com and ironport.com. You configure the entry for postmaster@ironport.com to bypass LDAP acceptance queries, and you configure the entry for ironport.com for ACCEPT. When you receive mail for postmaster@ironport.com, the LDAP acceptance bypass will occur only if the entry for postmaster@ironport.com is before the entry for ironport.com. If the entry for ironport.com is before the postmaster@ironport.com entry, the RAT matches the recipient address to this entry and applies the ACCEPT action.

Default RAT Entries

For all public listeners you create, by default, the RAT is set to *reject* email from *all* recipients:

ALL	REJECT

In the Recipient Access Table Overview listing, the default entry is named “All Other Recipients.”

Note By default, the RAT *rejects* all recipients so that you do not accidentally create an *open relay* on the Internet. An open relay (sometimes called an “insecure relay” or a “third-party” relay) is an SMTP email server that allows third-party relay of email messages. By processing mail that is neither for — nor from — a local user, an open relay makes it possible for an unscrupulous sender to route large volumes of spam through your gateway. Use caution when changing the default values of Recipient Access Tables for public listeners you create.

You can not delete the default “ALL” entry from the RAT.

Importing and Exporting Text Resources as Text Files

You will need access to the configuration directory on the appliance. Imported text files must be present in the configuration directory on the appliance. Exported text files are placed in the configuration directory.

See Appendix A, “Accessing the Appliance” for more information accessing on the configuration directory.

Modifying the RAT for a Listener via the GUI

To modify the RAT from the GUI, click Mail Policies > Recipient Access Table (RAT). The Recipient Access Table Overview page is displayed:

Figure 5-38 The Recipient Access Table Overview Page

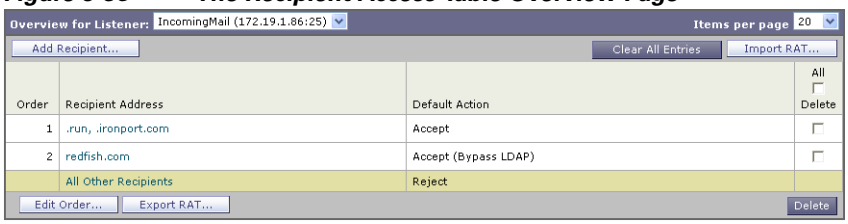

The Recipient Access Table Overview shows a listing of the entries in your RAT, including the order, default action, and whether or not the entry has been configured for bypassing LDAP accept queries.

From the Recipient Access Table Overview, you can:

- Add entries to the RAT
- Delete entries from the RAT
- Modify existing RAT entries
- Change the order of the entries
- Import RAT entries (overwrites existing entries) from a file
- Export RAT entries to a file

The RAT can be edited directly from the Command Line Interface (CLI). To customize a RAT for a listener you have defined, use the `edit -> rcptaccess -> new` subcommands of the `listenerconfig` command to add accepted local domains to the RAT for each public listener you configure. See the *Cisco IronPort AsyncOS CLI Reference Guide* for more information.

Adding New RAT Entries

Step 1 Click **Add Recipient** The Add to Recipient Access Table page is displayed:

Figure 5-39 Adding RAT Entries

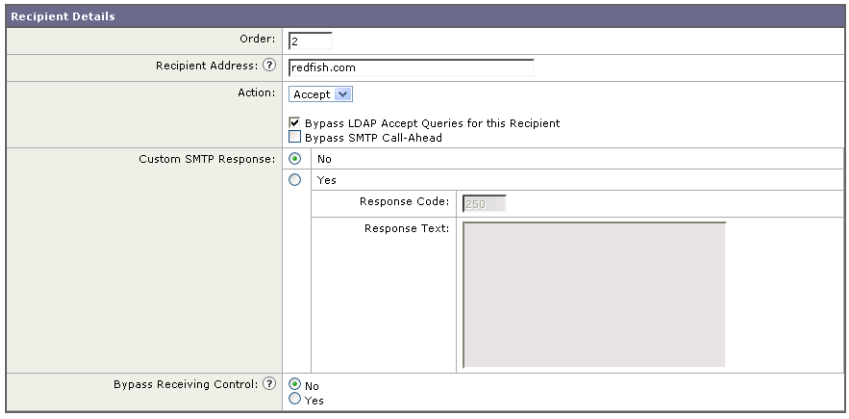

Step 2 Select an order for the entry.

Step 3 Enter the recipient address (see Defining Recipients, page 5-51 for more information about valid entries).

Step 4 Choose to accept or reject the recipient.

Step 5 Optionally, you can choose to bypass LDAP acceptance queries for the recipient (See Bypassing LDAP Accept for Special Recipients, page 5-52).

Step 6 If you want to use a custom SMTP response for this entry, select Yes for Custom SMTP Response. Enter a response code and text.

Step 7 Optionally, you can choose to bypass throttling (see Bypassing Throttling for Special Recipients, page 5-52) select Yes for Bypass Receiving Control.

Step 8 Submit and commit your changes.

Deleting RAT Entries

Step 1 Mark the checkbox in the Delete column for each entry you want to delete.

Step 2 Click **Delete**.

Step 3 The entry or entries you marked are removed from the RAT.

Step 4 Commit your changes.

Modifying RAT Entries

Step 1 Click the RAT entry in the Recipient Access Table Overview. The Edit Recipient Access Table page is displayed.

Step 2 Make changed to the entry.

Step 3 Commit your changes.

Changing the Order of RAT Entries

Step 1 Click **Edit Order** The Edit Recipient Access Table Order page is displayed:

Figure 5-40 Changing the Order of RAT Entries

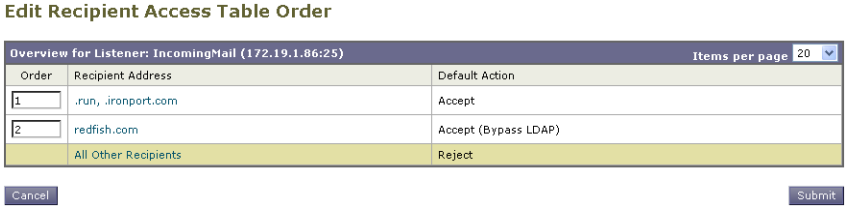

Step 2 Change the order by arranging the values in the Order column.

Step 3 Commit your changes.

Exporting RAT Entries

Step 1 Click **Export RAT** The Export Recipient Access Table page is displayed:

Figure 5-41 Exporting RAT Entries

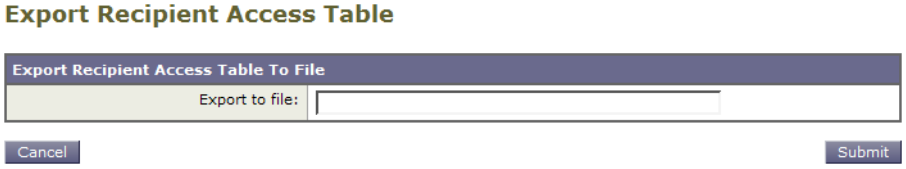

Step 2 Enter a file name for the exported entries. This is the name of the file that will be created in the configuration directory on the appliance.

Step 3 Submit and commit your changes.

Importing RAT Entries

When you import RAT entries, all of the existing RAT entries are removed from the RAT.

Step 1 Click **Import RAT** The Import Recipient Access Table page is displayed:

Figure 5-42 Exporting RAT Entries

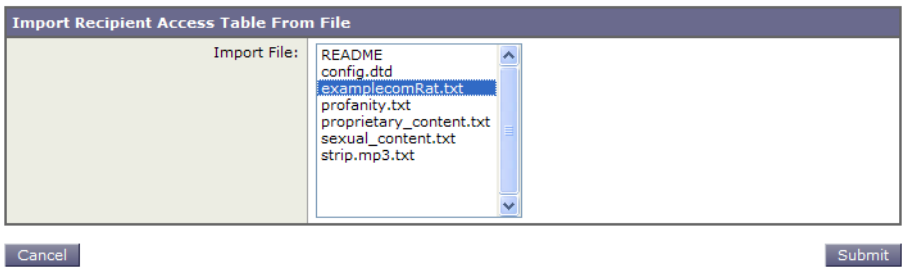

Step 2 Select a file from the list.

Note The file to import must be in the configuration directory on the appliance.

Step 3 Click **Submit**. You will see a warning message, asking you to confirm that you wish to remove all of the existing RAT entries.

Step 4 Click **Import**.

Step 5 Commit your changes.

You can place “comments” in the file. Lines that begin with a ‘#’ character are considered comments and are ignored by AsyncOS. For example:

```
# File exported by the GUI at 20060530T220526

.example.com  ACCEPT

ALL  REJECT
```

At this point, our Email Gateway configuration looks like this:

Figure 5-43 Editing the RAT for a Public Listener

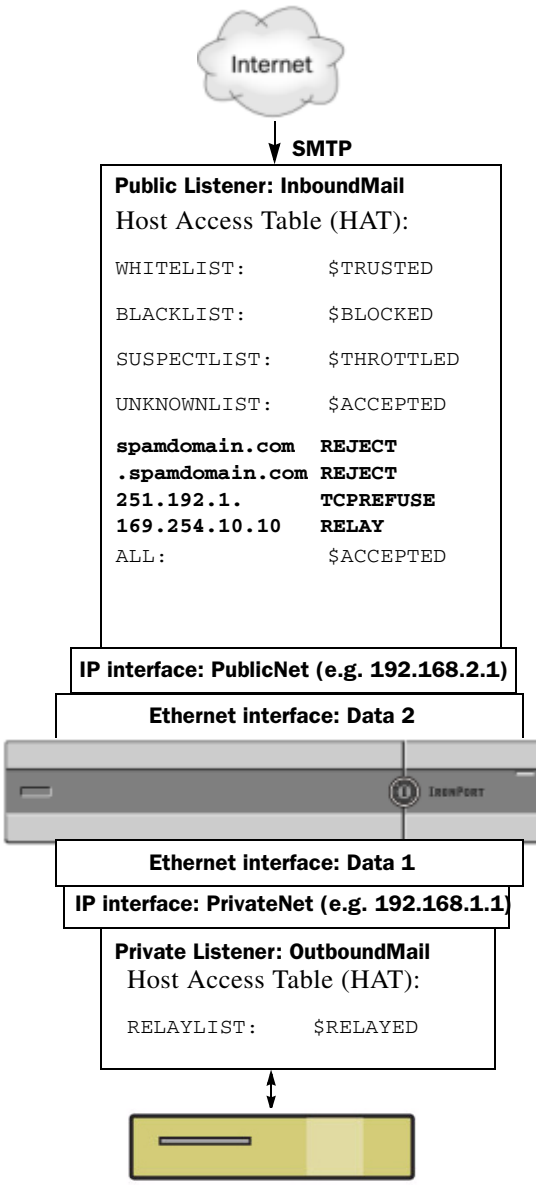

Note This public listener's RAT was modified to accept connections for the domain `newcompanyname.com`. A custom SMTP response was also created.

Note This private listener remains unchanged. Private listeners do not have a RAT.

Figure 5-44 expands the illustration shown in Figure 5-4 to include the processing sequence of a listener's HAT and (if applicable) RAT, and the default values for each.

Figure 5-44 Public and Private Listeners

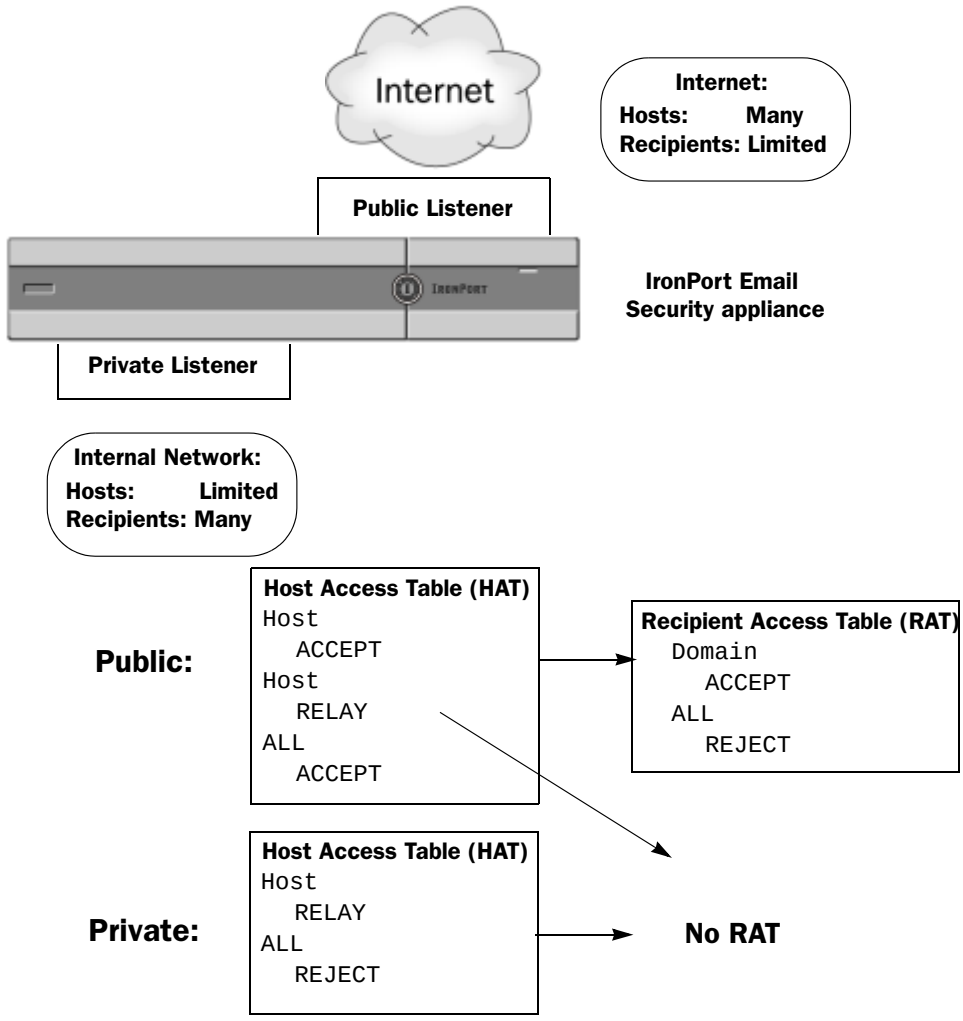

CHAPTER 6

Email Security Manager

Email Security Manager is a single, comprehensive dashboard to manage all email security services and applications on Cisco IronPort appliances. Prior to this release, the anti-spam and anti-virus settings were configured on a per-listener basis — meaning the policy was applied based on the receiving listener of an IP address, and not based on the recipient or sender of the message. Chapter 5, "Configuring the Gateway to Receive Email" describes how to create and configure listeners.

Email Security Manager allows you to manage the Outbreak Filters feature, anti-spam, anti-virus, and email content policies — on a per-recipient or per-sender basis, through distinct inbound and outbound policies.

Through the Mail Policies menu in the GUI (or the `policyconfig` command in the CLI), you create and manage incoming or outgoing mail policies. Mail policies are defined as a specific set of users (Envelope Recipients, Envelope Sender, From: header, or Reply-To: header) that map to specific settings for the following features:

- Anti-Spam Scanning
- Anti-Virus Scanning
- Outbreak Filters
- Content Filters
- RSA Email Data Loss Prevention Policies (outbound mail only)

Users can be defined by email address, email domains, or LDAP group queries.

Overview of User-Based Policies

User-based policies in Email Security Manager are designed to allow you to create the policies that satisfy the different and sometimes disparate security needs of all users within your organization.

For example, using this feature, you can quickly create policies to enforce the following conditions:

- Disable Cisco IronPort Anti-Spam scanning for all email to the Sales organization. Enable it for the Engineering organization with a moderate policy: tag the subject lines of suspected spam and legitimate marketing messages, and drop positively identified spam. For the Human Resources organization, enable anti-spam scanning with an aggressive policy: quarantine suspected spam messages, quarantine legitimate marketing messages, and drop positively identified spam.

- Drop dangerous executable attachments for all users except those in the System Administrator group.
- Scan and attempt to repair viruses in messages destined for the Engineering organization, but drop infected attachments for all messages sent to the address `jobs@example.com`.
- Scan all outgoing messages using RSA Email Data Loss Prevention (DLP) for possible confidential information. If a message matches, quarantine the message and send a blind-carbon copy to the Legal department.

Note If you are using RSA Enterprise Manager for DLP, the outgoing mail policy is assigned to a DLP policy in Enterprise Manager. See RSA Enterprise Manager, page 11-27 for more information.

- If an incoming message contains an MP3 attachment, quarantine the message and send a message to the intended recipient with instructions for calling the Network Operations Center to retrieve the message. Expire such messages after 10 days.
- Include a disclaimer to all outgoing mail from the Executive Staff with the company's newest tag line, but include a different "forward-looking statements" disclaimer to all outgoing mail from the Public Relations organization.
- Enable the Outbreak Filters feature for all incoming messages, but bypass scanning for messages with links to example.com or attachments whose file extension is `.dwg`.

Note Content dictionaries, disclaimers, and notification templates must be created before they can be referenced by content filters. For more information, see Text Resources, page 14-1.

Incoming vs. Outgoing Messages

Two policy tables are defined in the Email Security Manager: one table for messages from sending hosts that are stipulated by HAT policies with the "Accept" behavior, the other table for sending hosts qualified as having HAT "Relay" behavior. The former table is the *incoming* policy table, the latter is the *outgoing* policy table.

- *Incoming messages* are messages received from connections that match an ACCEPT HAT policy in any listener.
- *Outgoing messages* are messages from connections that match a RELAY HAT policy in any listener. This includes any connection that was authenticated with SMTP AUTH.

Note In certain installations, "internal" mail being routed through the Cisco IronPort appliance will be considered *outgoing*, even if all the recipients are addressed to internal addresses. For example, by default for Cisco IronPort C10/100 customers, the system setup wizard will configure only one physical Ethernet port with one listener for receiving inbound email and relaying outbound email.

For many configurations, you can think of the incoming table as Public, while the Outgoing table is Private, although both could be used by a single listener. The policy table used on a particular message is not dependant on the direction of the message, with respect to sender or recipient addresses, out to the internet or in to an intranet.

You manage these tables using the Mail Policies > Incoming Mail Policies or Outgoing Mail Policies pages in the GUI, or the `policyconfig` command in the CLI. You can assign individual mail policies to delegated administrators whose responsibilities include managing your mail system. See the "Common Administrative Tasks" chapter in *Cisco IronPort AsyncOS for Email Daily Management Guide* for more information.

Note DLP scanning can only be performed on outgoing messages.

Policy Matching

As incoming messages are received by listeners on the system, each message recipient matches a policy in one of the tables, regardless of the number of listeners configured on the system. Matches are based on either the recipient's address or the sender's address:

- Recipient address matches the Envelope Recipient address

  When matching recipient addresses, the recipient addresses entered are the final addresses after processing by preceding parts of the email pipeline. For example, if enabled, the default domain, LDAP routing or masquerading, alias table, domain map, and message filters features can rewrite the Envelope Recipient address and may affect whether the message matches a policy in the Email Security Manager (Anti-Spam, Anti-Virus, Content Filters, and Outbreak Filters).

- Sender address matches:
  - Envelope Sender (RFC821 MAIL FROM address)
  - Address found in the RFC822 From: header
  - Address found in the RFC822 Reply-To: header

Addresses may be matched on either a full email address, user, domain, or partial domain, and addresses may also match LDAP group membership.

First Match Wins

Each recipient is evaluated for each policy in the appropriate table (incoming or outgoing) in a top-down fashion.

For each recipient of a message, the first matching policy wins. If a recipient does not match any specific policy, the recipient will automatically match the default policy of the table.

If a match is made based on a sender address (or on the special "Listener" rule created by an upgrade — see below), all remaining recipients of a message will match that policy. (This is because there can be only one sender or one listener per message.)

Examples of Policy Matching

The following examples help show how the policy tables are matched in a top-down fashion.

Given the following Incoming Mail Email Security Policy table shown in Table 6-1, incoming messages will match different policies.

Table 6-1 Policy Matching Example

Order	Policy Name	Users
1	special_people	Recipient: joe@example.com Recipient: ann@example.com
2	from_lawyers	Sender: @lawfirm.com
3	acquired_domains	Recipient: @newdomain.com Recipient: @anotherexample.com
4	engineering	Recipient: PublicLDAP.ldapgroup: engineers
5	sales_team	Recipient: jim@ Recipient: john@ Recipient: larry@
	Default Policy	(all users)

Example 1

A message from sender bill@lawfirm.com sent to recipient jim@example.com will match policy #2, because the user description that matches the sender (@lawfirm.com) appears sooner in the table than the user description that matches the recipient (jim@).

Example 2

Sender joe@yahoo.com sends an incoming message with three recipients: john@example.com, jane@newdomain.com, and bill@example.com. The message for recipient jane@newdomain.com will receive the anti-spam, anti-virus, outbreak filters, and content filters defined in policy #3, while the message for recipient john@example.com will receive the settings defined in policy #5. Because the recipient bill@example.com does not match the engineering LDAP query, the message will receive the settings defined by the default policy. This example shows how messages with multiple recipients can incur *message splintering*. See Message Splintering, page 6-4 for more information.

Example 3

Sender bill@lawfirm.com sends a message to recipients ann@example.com and larry@example.com. The recipient ann@example.com will receive the anti-spam, anti-virus, outbreak filters, and content filters defined in policy #1, and the recipient larry@example.com will receive the anti-spam, anti-virus, outbreak filters, and content filters defined in policy #2, because the sender (@lawfirm.com) appears sooner in the table than the user description that matches the recipient (jim@).

Message Splintering

Intelligent message splintering (by matching policy) is the mechanism that allows for differing recipient-based policies to be applied independently to message with multiple recipients.

Each recipient is evaluated for each policy in the appropriate Email Security Manager table (incoming or outgoing) in a top-down fashion.

Each policy that matches a message creates a new message with those recipients. This process is defined as *message splintering*:

- If some recipients match different policies, the recipients are grouped according to the policies they matched, the message is split into a number of messages equal to the number of policies that matched, and the recipients are set to each appropriate "splinter."
- If all recipients match the same policy, the message is not splintered. Conversely, a maximum splintering scenario would be one in which a single message is splintered for each message recipient.
- Each message splinter is then processed by anti-spam, anti-virus, DLP scanning (outgoing messages only), Outbreak Filters, and content filters independently in the email pipeline.

Table 6-2 illustrates the point at which messages are splintered in the email pipeline.

Note Email DLP scanning is only available for outgoing messages.

Table 6-2 Message Splintering in the Email Pipeline

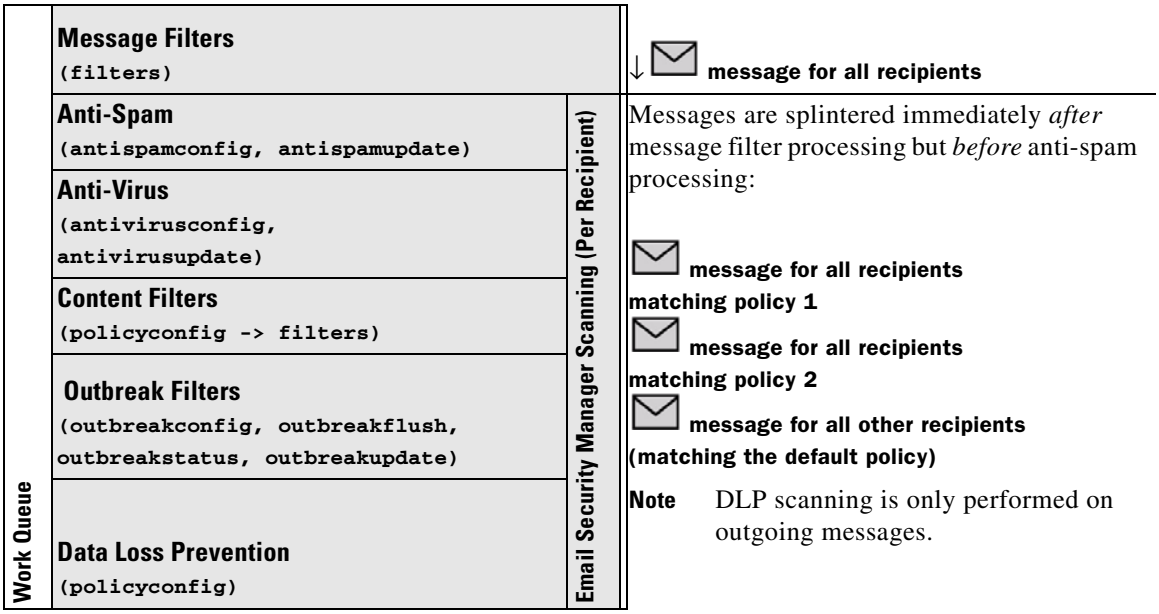

Note New MIDs (message IDs) are created for each message splinter (for example, MID 1 becomes MID 2 and MID 3). For more information, see the "Logging" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*. In addition, the trace function shows which policies cause a message to be split.

Policy matching and message splintering in Email Security Manager policies obviously affect how you manage the message processing available on the appliance.

Managed Exceptions

Because the iterative processing of each splinter message impacts performance, Cisco recommends using the Incoming and Outgoing Mail Policies tables of Email Security Manager to configure policies on a *managed exception* basis. In other words, evaluate your organization's needs and try to configure the feature so that the majority of messages will be handled by the default policy and the minority of

messages will be handled by a few additional "exception" policies. In this manner, message splintering will be minimized and you are less likely to impact system performance from the processing of each splinter message in the work queue.

Contents of Policies

Email Security Manager tables match incoming or outgoing messages for specific groups of users (Envelope Recipients, Envelope Sender, From: header, or Reply-To: header) and map them to specific settings for the following features:

- Anti-Spam Scanning — See Anti-Spam, page 9-1 for more information.
- Anti-Virus Scanning — See Anti-Virus, page 8-1 for more information.
- Content Filters — See Content Filters Overview, page 6-6 for more information.
- Outbreak Filters

  Cisco IronPort's Outbreak Filters feature is a predictive security service that provides a "first line of defense" against new virus, phishing, and scam outbreaks by quarantining suspicious messages until traditional anti-virus and anti-spam security services can be updated to detect them. You can enable or disable Outbreak filters for given recipients, and also define the file types that will bypass the Outbreak Filters feature in Email Security Manager. See Chapter 10, "Outbreak Filters" for more information.
- Data Loss Prevention — See Chapter 11, "Data Loss Prevention" for more information.

Figure 6-1 illustrates the Email Security Manager in the GUI that maps users defined in a policy to specific Anti-Spam, Anti-Virus, Outbreak Filter, DLP, and Content Filters settings.

Figure 6-1 Summary of Email Security Manager Policies in the GUI

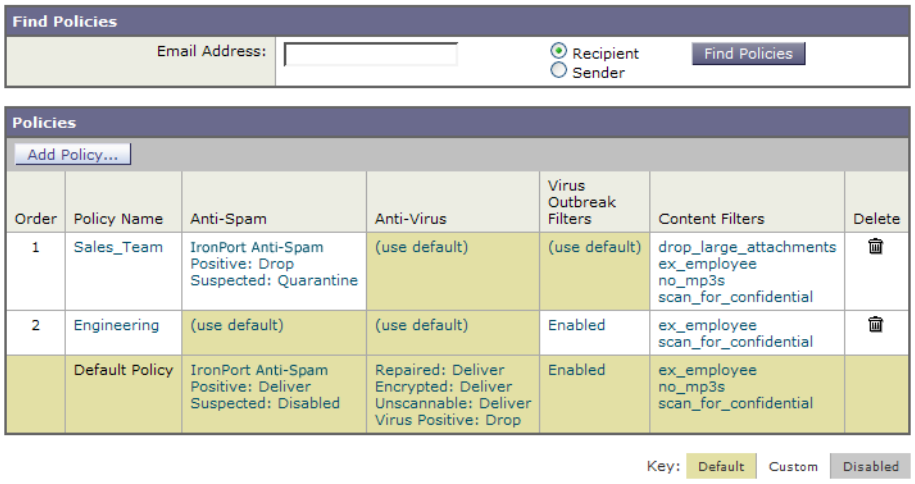

Content Filters Overview

Email Security Manager policies allow you to create content filters to be applied to messages on a per-recipient or per-sender basis. Content filters are similar to message filters, except that they are applied later in the email pipeline — after a message has been "splintered" into a number of separate

messages for each matching Email Security Manager policy. The functionality of content filters is applied after message filters processing and anti-spam and anti-virus scanning have been performed on a message.

Like regular message filters, you define a name for each content filter. The name must be unique to the Incoming or Outgoing Mail Policies table in which it will be used. Each Incoming and Outgoing Mail Policies table will have its own, singular "master list" of content filters. The order is defined on a per-table basis (for incoming or outgoing). However, each individual policy determines which particular filters will be executed.

Unlike regular message filters (which are applied before anti-spam and anti-virus scanning), content filters can be configured both in the CLI and in the GUI. The GUI includes a "rule builder" page that allows you to easily create the conditions and actions that constitute a content filter. Email Security Manager incoming or outgoing mail policy tables manage which content filters are enabled the order in which they will be applied for any given policy. Table 6-3 lists the available *conditions* you can use to create a content filter. Table 6-4 lists the non-final and final *actions* you can use to define a content filter. Together, conditions and action constitute a content filter. Table 6-5 shows the action variables you can use when creating content filters.

You can specify which delegated administration user roles can edit the content filter and enable them in mail policies. For more information on delegated administrators' access privileges for content filters, see the "Common Administrative Tasks" chapter in *Cisco IronPort AsyncOS for Email Daily Management Guide*.

Content Filter Conditions

Specifying conditions in content filters is optional.

In the content filter conditions, when you add filter rules that search for patterns in the message body or attachments, you can specify the minimum threshold for the number of times the pattern must be found. When AsyncOS scans the message, it totals the "score" for the number of matches it finds in the message and attachments. If the minimum threshold is not met, the regular expression does not evaluate to true. You can specify this threshold for text, smart identifiers, or content dictionary terms.

You can also use "smart identifiers" to identify patterns in data. Smart identifiers can detect the following patterns:

- Credit card numbers
- U.S. Social Security numbers
- CUSIP (Committee on Uniform Security Identification Procedures) numbers
- ABA (American Banking Association) routing numbers

For more information about specifying a minimum threshold for the number of times a pattern must be found, and smart identifiers, see the "Using Message Filters to Enforce Email Policies" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Multiple conditions may defined for each filter. When multiple conditions are defined, you can choose whether the conditions are tied together as a logical OR ("Any of the following conditions...") or a logical AND ("All of the following conditions").

Table 6-3 Content Filter Conditions

Condition	Description
(no conditions)	Specifying conditions in content filters is optional. If no conditions are specified, a `true` rule is implied. The `true` rule matches all messages, and the actions are always performed.
Message Body or Attachments	**Contains text:** Does the message body contain text or an attachment that matches a specific pattern? **Contains smart identifier:** Does content in the message body or attachment match a smart identifier? **Contains term in content dictionary:** Does the message body contain any of the regular expressions or terms in the content dictionary named <*dictionary name*>? For this option to be enabled, the dictionary must already have been created. See Content Dictionaries, page 14-2. **Number of matches required**. Specify the number of matches required for the rule to evaluate to true. You can specify this threshold for text, smart identifiers, or content dictionary terms. This includes delivery-status parts and associated attachments.
Message Body	**Contains text:** Does the message body contain text that matches a specific pattern? **Contains smart identifier:** Does content in the message body match a smart identifier? **Contains term in content dictionary:** Does the message body contain any of the regular expressions or terms in the content dictionary named <*dictionary name*>? For this option to be enabled, the dictionary must already have been created. See Content Dictionaries, page 14-2. **Number of matches required**. Specify the number of matches required for the rule to evaluate to true. You can specify this threshold for text or smart identifiers. This rule applies to the body of the message only. It does not include attachments or headers.

Table 6-3 Content Filter Conditions (Continued)

Condition	Description
Message Size	Is the body size within a specified range? Body size refers to the size of the message, including both headers and attachments. The body-size rule selects those messages where the body size compares as directed to a specified number.
Attachment Content	**Contains text.** Does the message contain an attachment that contains text or another attachment that matches a specific pattern? This rule is similar to the `body-contains()` rule, but it attempts to avoid scanning the entire "body" of the message. That is, it attempts to scan only that which the user would view as being an attachment. **Contains a smart identifier.** Does content in the message attachment match the specified smart identifier? **Contains terms in content dictionary.** Does the attachment contain any of the regular expressions or terms in the content dictionary named *<dictionary name>*? To search for dictionary terms, the dictionary must already have been created. See Content Dictionaries, page 14-2. **Number of matches required**. Specify the number of matches required for the rule to evaluate to true. You can specify this threshold for text, smart identifier, or content dictionary matches.
Attachment File Info	**Filename.** Does the message contain an attachment with a filename that matches a specific pattern? **File type**. Does the message contain an attachment of a file type that matches a specific pattern based on its fingerprint (similar to a UNIX `file` command)? **MIME type**. Does the message contain an attachment of a specific MIME type? This rule is similar to the `attachment-type` rule, except only the MIME type given by the MIME attachment is evaluated. (The appliance does not try to "guess" the type of the file by its extension if there is no explicit type given.) **Image Analysis.** Does the message contain an image attachment that matches the image verdict specified? Valid image analysis verdicts include: *Suspect, Inappropriate, Suspect or Inappropriate, Unscannable*, or *Clean.*

Table 6-3 Content Filter Conditions (Continued)

Condition	Description
Attachment Protection	**Contains an attachment that is password-protected or encrypted.** (For example, use this condition to identify attachments that are potentially unscannable.) **Contains an attachment that is NOT password-protected or encrypted.**
Subject Header	**Subject Header:** Does the subject header match a certain pattern? **Contains terms in content dictionary**: Does the subject header contain any of the regular expressions or terms in the content dictionary <*dictionary name*>? To search for dictionary terms, the dictionary must already have been created. See Content Dictionaries, page 14-2.
Other Header	**Header name**: Does the message contain a specific header? **Header value:** Does the value of that header match a certain pattern? **Header value contains terms in the content dictionary.** Does the specified header contain any of the regular expressions or terms in the content dictionary named <*dictionary name*>? To search for dictionary terms, the dictionary must already have been created. See Content Dictionaries, page 14-2
Envelope Sender	**Envelope Sender.** Does the Envelope Sender (i.e., the Envelope From, <MAIL FROM>) match a given pattern? **Matches LDAP group**. Is the Envelope Sender, i.e., the Envelope From, <MAIL FROM>) in a given LDAP group? **Contains term in content dictionary.** Does the envelope sender contain any of the regular expressions or terms in the content dictionary named <*dictionary name*>? To search for dictionary terms, the dictionary must already have been created. See Content Dictionaries, page 14-2.

Table 6-3 Content Filter Conditions (Continued)

Condition	Description
Envelope Recipient	**Envelope Recipient**. Does the Envelope Recipient, (i.e. the Envelope To, <RCPT TO>) match a given pattern? **Matches LDAP group**. Is the Envelope Recipient, (i.e. the Envelope To, <RCPT TO>) in a given LDAP group? **Contains term in content dictionary.** Does the envelope recipient contain any of the regular expressions or terms in the content dictionary named <*dictionary name*>? To search for dictionary terms, the dictionary must already have been created. See Content Dictionaries, page 14-2. **Note:** The Envelope Recipient rule is message-based. If a message has multiple recipients, only one recipient has to be found in a group for the specified action to affect the message to all recipients. Is the Envelope Sender (i.e., the Envelope From, <MAIL FROM>) in a given LDAP group?
Receiving Listener	Did the message arrive via the named listener? The listener name must be the name of a listener currently configured on the system.
Remote IP	Was the message sent from a remote host that matches a given IP address or IP block? The Remote IP rule tests to see if the IP address of the host that sent that message matches a certain pattern. This can be an Internet Protocol version 4 (IPv4) or version 6 (IPv6) address. The IP address pattern is specified using the allowed hosts notation described in Sender Group Syntax, page 5-21, except for the SBO, SBRS, dnslist notations and the special keyword ALL.
Reputation Score	What is the sender's SenderBase Reputation Score? The Reputation Score rule checks the SenderBase Reputation Score against another value.
DKIM Authentication	Did DKIM authentication pass, partially verify, return temporarily unverifiable, permanently fail, or were no DKIM results returned?
SPF Verification	What was the SPF verification status? This filter rule allows you to query for different SPF verification results. For more information about SPF verification, see "Email Authentication" in *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Note The dictionary-related conditions are only available if you have one or more dictionaries enabled. For information about creating content dictionaries, see Content Dictionaries, page 14-2.

Figure 6-2 Content Filter Conditions

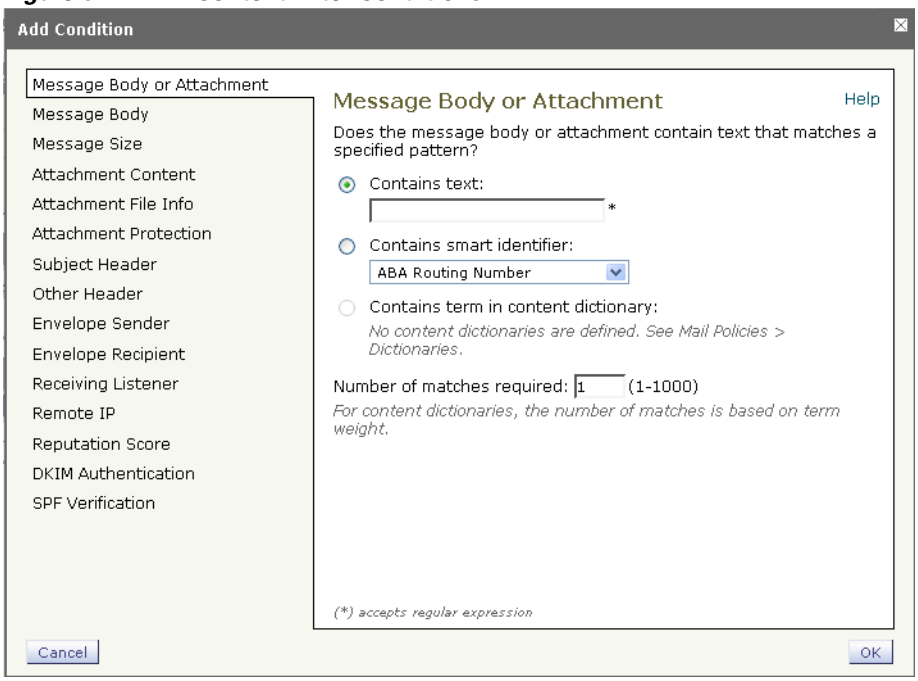

Content Filter Actions

At least one action must be defined for each content filter.

Actions are performed in order on messages, so consider the order of actions when defining multiple actions for a content filter.

When you configure a quarantine action for messages that match Attachment Content conditions, Message Body or Attachment conditions, Message body conditions, or the Attachment content conditions, you can view the matched content in the quarantined message. When you display the message body, the matched content is highlighted in yellow. You can also use the $MatchedContent action variable to include the matched content in the message subject. For more information, see *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Only one final action may be defined per filter, and the final action must be last action listed. Bounce, deliver, and drop are final actions. When entering actions for content filters, the GUI and CLI will force final actions to be placed last.

Table 6-4 Content Filter Actions

Action	Description
Quarantine	**Quarantine**. Flags the message to be held in one of the system quarantine areas. **Duplicate message**: Sends a copy of the message to the specified quarantine and continues processing the original message. Any additional actions apply to the original message.
Encrypt on Delivery	The message continues to the next stage of processing. When all processing is complete, the message is encrypted and delivered. **Encryption rule**: Always encrypts the message or only encrypts it if an attempt to send it over a TLS connection first fails. See Using a TLS Connection as an Alternative to Encryption, page 12-8 for more information. **Encryption Profile.** Once processing is complete, encrypts the message using the specified encryption profile, then delivers the message. This action is for use with a Cisco IronPort Encryption Appliance or a hosted key service. **Subject**. Subject for the encrypted message. By default, the value is `$Subject`.
Strip Attachment by Content	**Attachment contains.** Drops all attachments on messages that contain the regular expression. Archive files (zip, tar) will be dropped if any of the files they contain match the regular expression pattern. **Contains smart identifier.** Drops all attachments on a message that contains the specified smart identifier. **Attachment contains terms in the content dictionary.** Does the attachment contain any of the regular expressions or terms in the content dictionary named <*dictionary name*>? **Number of matches required**. Specify the number of matches required for the rule to evaluate to true. You can specify this threshold for text, smart identifier, or content dictionary matches. **Replacement message.** The optional comment serves as the means to modify the text used to replace the attachment that was dropped. Attachment footers simply append to the message.

Table 6-4 Content Filter Actions (Continued)

Action	Description
Strip Attachment by File Info	**File name**. Drops all attachments on messages that have a filename that match the given regular expression. Archive file attachments (zip, tar) will be dropped if they contain a file that matches. **File size**. Drops all attachments on the message that, in raw encoded form, are equal to or greater than the size (in bytes) given. Note that for archive or compressed files, this action does not examine the uncompressed size, but rather the size of the actual attachment itself. **File type**. Drops all attachments on messages that match the given "fingerprint" of the file. Archive file attachments (zip, tar) will be dropped if they contain a file that matches. **MIME type**. Drops all attachments on messages that have a given MIME type. **Image Analysis Verdict**. Drops attachments for image attachments that match the image verdict specified. Valid image analysis verdicts include: *Suspect, Inappropriate, Suspect or Inappropriate, Unscannable*, or *Clean*. **Replacement message.** The optional comment serves as the means to modify the text used to replace the attachment that was dropped. Attachment footers simply append to the message.
Add Disclaimer Text	**Above**. Add disclaimer above message (heading). **Below.** Add disclaimer below message (footer). **Note:** You must have already created disclaimer text in order to use this content filter action. See Disclaimer Template, page 14-17 for more information.
Bypass Outbreak Filter Scanning	Bypass Outbreak Filter scanning for this message.
Bypass DKIM Signing	Bypass DKIM signing for this message.

Table 6-4 Content Filter Actions (Continued)

Action	Description
Send Copy (Bcc:)	**Email addresses**. Copies the message anonymously to the specified recipients. **Subject**. Add a subject for the copied message. **Return path (optional).** Specify a return path. **Alternate mail host (optional).** Specify an alternate mail host.
Notify	**Notify.** Reports this message to the specified recipients. You can optionally notify the sender and recipients. **Subject**. Add a subject for the copied message. **Return path (optional).** Specify a return path. **Use template**. Select a template from the templates you created. **Include original message as an attachment**. Adds the original message as an attachment.
Change Recipient to	**Email address**. Changes the recipient of the message to the specified email address.
Send to Alternate Destination Host	**Mail host**. Changes the destination mail host for the message to the specified mail host. **Note** This action prevents a message classified as spam by an anti-spam scanning engine from being quarantined. This action overrides the quarantine and sends it to the specified mail host.
Deliver from IP Interface	**Send from IP interface.** Send from the specified IP Interface. The Deliver from IP Interface action changes the source host for the message to the source specified. The source host consists of the IP interface that the messages should be delivered from.
Strip Header	**Header name.** Remove the specified header from the message before delivering.

Table 6-4 Content Filter Actions (Continued)

Action	Description
Add/Edit Header	**Inserts a new header into the message or modifies an existing header.** **Header name**. Name of new or existing header. **Specify value of new header.** Inserts a value for the new header into the message before delivering. **Prepend to the Value of Existing Header.** Prepends the value to the existing header before delivering. **Append to the Value of Existing Header.** Appends the value to the existing header before delivering. **Search & Replace from the Value of Existing Header**. Enter a search term to find the value you want to replace in the existing header in the **Search for** field. Enter the value you want to insert into the header in the **Replace with** field. You can use a regular expression to search for the value. Leave the **Replace with** field empty if you want to delete the value from the header.
Add Message Tag	Inserts a custom term into the message to use with RSA Email DLP policy filtering. You can configure a RSA Email DLP policy to limit scanning to messages with the message tag. The message tag is not visible to recipients. For information on using messages tags in a DLP policy, see DLP Policies, page 11-10.
Add Log Entry	Inserts customized text into the IronPort Text Mail logs at the `INFO` level. The text can include action variables. The log entry also appears in message tracking.
Encrypt and Deliver Now (Final Action)	Encrypts and delivers the message, skipping any further processing. **Encryption rule**: Always encrypts the message or only encrypts it if an attempt to send it over a TLS connection first fails. See Using a TLS Connection as an Alternative to Encryption, page 12-8 for more information. **Encryption Profile.** Encrypts the message using the specified encryption profile, then delivers the message. This action is for use with a Cisco IronPort Encryption Appliance or a hosted key service. **Subject**. Subject for the encrypted message. By default, the value is `$Subject`.
Bounce (Final Action)	Sends the message back to the sender.
Skip Remaining Content Filters (Final Action)	Delivers the message to the next stage of processing, skipping any further content filters. Depending on configuration, this may mean deliver the message to recipient(s), quarantine, or begin Outbreak Filters scanning.
Drop (Final Action)	Drops and discards the message.

Figure 6-3 Content Filter Actions in GUI

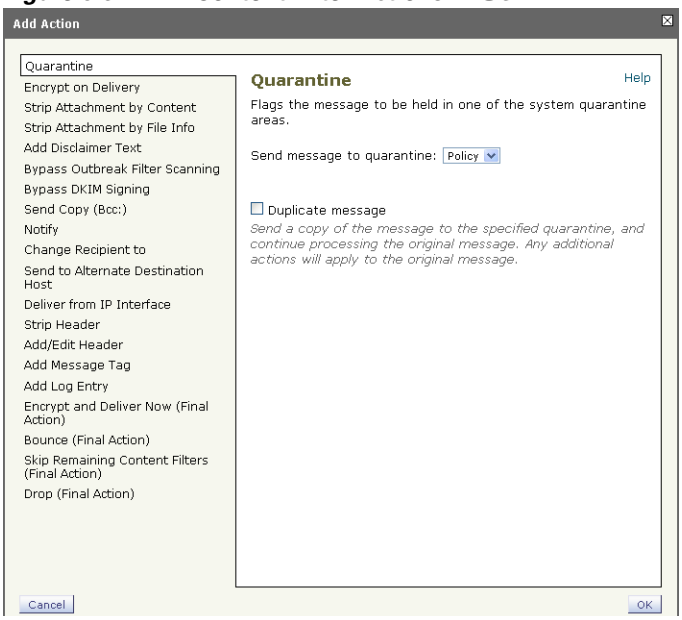

Action Variables

Headers added to messages processed by content filters can contain variables that will be automatically replaced with information from the original message when the action is executed. These special variables are called *action variables.* Your Cisco IronPort appliance supports the following set of action variables:

Table 6-5 Action Variables

Variable	Syntax	Description
All Headers	`$AllHeaders`	Replaced by the message headers.
Body Size	`$BodySize`	Replaced by the size, in bytes, of the message.
Date	`$Date`	Replaced by the current date, using the format MM/DD/YYYY.
Dropped File Name	`$dropped_filename`	Returns only the most recently dropped filename.
Dropped File Names	`$dropped_filenames`	Same as `$filenames`, but displays list of dropped files.
Dropped File Types	`$dropped_filetypes`	Same as `$filetypes`, but displays list of dropped file types.
Envelope Sender	`$envelopefrom` or `$envelopesender`	Replaced by the Envelope Sender (Envelope From, <MAIL FROM>) of the message.
Envelope Recipients	`$EnvelopeRecipients`	Replaced by all Envelope Recipients (Envelope To, <RCPT TO>) of the message.
File Names	`$filenames`	Replaced with a comma-separated list of the message's attachments' filenames.
File Sizes	`$filesizes`	Replaced with a comma-separated list of the message's attachment's file sizes.

Table 6-5 Action Variables (Continued)

Variable	Syntax	Description
File Types	`$filetypes`	Replaced with a comma-separated list of the message's attachments' file types.
Filter Name	`$FilterName`	Replaced by the name of the filter being processed.
GMTimeStamp	`$GMTimeStamp`	Replaced by the current time and date, as would be found in the Received: line of an email message, using GMT.
HAT Group Name	`$Group`	Replaced by the name of the sender group the sender matched on when injecting the message. If the sender group had no name, the string "`>Unknown<`" is inserted.
Mail Flow Policy	`$Policy`	Replaced by the name of the HAT policy applied to the sender when injecting the message. If no predefined policy name was used, the string "`>Unknown<`" is inserted.
Matched Content	`$MatchedContent`	Replaced by the value (or values) that triggered a content-scanning filter. Matched content can be a content dictionary match, a smart identifier, or a match to a regular expression.
Header	`$Header['`*`string`*`']`	Replaced by the value of the quoted header, if the original message contains a matching header. Note that double quotes may also be used.
Hostname	`$Hostname`	Replaced by the hostname of the Cisco IronPort appliance.
Internal Message ID	`$MID`	Replaced by the Message ID, or "MID" used internally to identify the message. Not to be confused with the RFC822 "Message-Id" value (use $Header to retrieve that).
Receiving Listener	`$RecvListener`	Replaced by the nickname of the listener that received the message.
Receiving Interface	`$RecvInt`	Replaced by the nickname of the interface that received the message.
Remote IP Address	`$RemoteIP`	Replaced by the IP address of the system that sent the message to the Cisco IronPort appliance.
Remote Host Address	`$remotehost`	Replaced by the hostname of the system that sent the message to the Cisco IronPort appliance.
SenderBase Reputation Score	`$Reputation`	Replaced by the SenderBase Reputation score of the sender. If there is no reputation score, it is replaced with "`None`".
Subject	`$Subject`	Replaced by the subject of the message.

Table 6-5 Action Variables (Continued)

Variable	Syntax	Description
Time	`$Time`	Replaced by the current time, in the local time zone.
Timestamp	`$Timestamp`	Replaced by the current time and date, as would be found in the Received: line of an email message, in the local time zone.

Practical Example (GUI)

The following example demonstrates the features of Email Security Manager by illustrating the following tasks:

Step 1 Editing the anti-spam, anti-virus, Outbreak Filter, and Content Filters for the default Incoming Mail Policy.

Step 2 Adding two new policies for different sets of users — the sales organization and the engineering organization — and then configuring different email security settings for each.

Step 3 Creating three new content filters to be used in the Incoming Mail Overview policy table.

Step 4 Editing the policies again to enable the content filters for some groups, but not for others.

This example is meant to show the power and flexibility with which you can manage different recipient-based settings for anti-spam, anti-virus, Outbreak Filter, and Content Filters in Email Security Manager. This example assigns these a custom user role called “Policy Administrator” that has mail policy and content filters access privileges. For more detailed information about how anti-spam, anti-virus, Outbreak filters, and delegated administration work, refer to the chapters following this one:

- Anti-Spam, page 9-1
- Anti-Virus, page 8-1
- Outbreak Filters, page 10-1
- Common Administrative Tasks, *Cisco IronPort AsyncOS for Email Daily Management Guide*

Accessing Email Security Manager

On newly-installed or upgraded systems, access Email Security Manager by clicking the Mail Policies tab. By default, The Incoming Mail Policies table is displayed.

On brand new systems, if you completed all steps in the system setup wizard and you chose to enable Cisco IronPort Anti-Spam, Sophos or McAfee Anti-Virus, and Outbreak Filters, the Incoming Mail Policies Page will resemble Figure 6-4.

By default, these settings are enabled for the default Incoming Mail Policy:

- Anti-Spam (if the Cisco IronPort Spam Quarantine is enabled): Enabled
  - Positively-identified spam: quarantine, prepend the message subject
  - Suspected spam: quarantine, prepend the message subject
  - Marketing email: scanning not enabled
- Anti-Spam (if the Cisco IronPort Spam Quarantine is not enabled): Enabled

  - Positively-identified spam: deliver, prepend the message subject
  - Suspected spam: deliver, prepend the message subject
  - Marketing email: scanning not enabled
- Anti-Virus: Enabled, Scan and Repair viruses, include an X-header with anti-virus scanning results
  - Repaired messages: deliver, prepend the message subject
  - Encrypted messages: deliver, prepend the message subject
  - Unscannable messages: deliver, prepend the message subject
  - Virus infected messages: drop
- Outbreak Filters: Enabled
  - No file extensions are excepted
  - Retention time for messages with suspect viral attachments is 1 day
  - Message modification is not enabled
- Content Filters: Disable

Figure 6-4 Incoming Mail Policies Page: Defaults for a Brand New Appliance

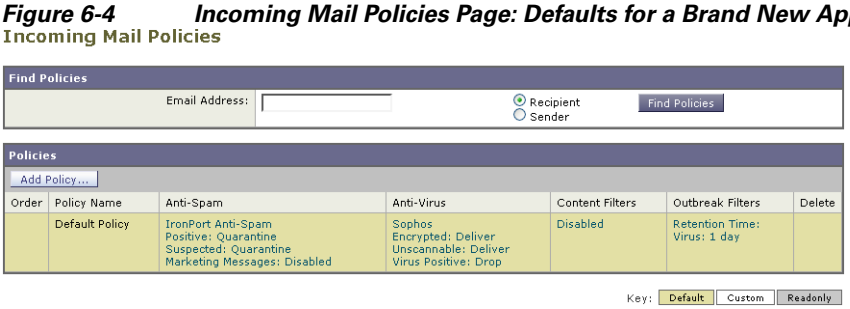

Note In this example, the Incoming Mail Policy will use the default anti-spam settings for when the Cisco IronPort Spam Quarantine is enabled.

Enabled, Disabled, and "Not Available"

The columns in an Email Security Manager table (either incoming or outgoing) display links for the state of the security service for each policy name. If a service is enabled, the word "Enabled" or a summary of the configuration is displayed. Similarly, the word "Disabled" is displayed if a service is disabled.

"Not Available" is displayed as a link if the license agreement for a service has not been accepted yet or a service has expired. In these cases, clicking the "Not Available" link will display the global page within the Security Services tab, rather than the page where you can configure per-policy settings for a service. An alert is displayed to let you know that your page has changed to a different tab. See Figure 6-5.

Figure 6-5 Security Services Not Available

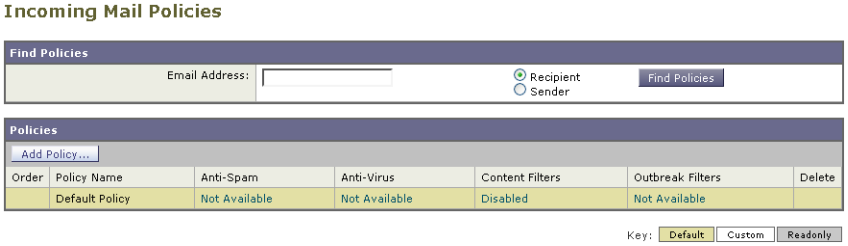

Editing the Default Policy: Anti-Spam Settings

Each row in the Email Security Manager represents a different policy. Each column represents a different security service.

- To edit the default policy, click any of the links for a security service in the bottom row of the Email Security Manager incoming or outgoing mail policy table.

In this example, you will change the anti-spam settings for the default policy for incoming mail to be more aggressive. The default value is to quarantine positively identified and suspected spam messages, with marketing email scanning disabled. This example shows how to change the setting so that positively identified spam is dropped. Suspected spam continues to be quarantined. Marketing email scanning is enabled, with marketing messages being delivered to the intended recipients. The subjects of marketing messages will be prepended with the text `[MARKETING]`.

Step 1 Click the link for the anti-spam security service. The Anti-Spam settings page shown in Figure 6-6 is displayed.

Note For default security service settings, the first setting on the page defines whether the service is enabled for the policy. You can click “Disable” to disable the service altogether.

Step 2 In the “Positively Identified Spam Settings” section, change the “Action to apply to this message” to Drop.

Step 3 In the “Marketing Email Settings” section, click **Yes** to enable marketing email scanning.

If enabled, the default action is to deliver legitimate marketing messages while prepending the subject with the text `[MARKETING]`.

The “Add text to message” field only accepts US-ASCII characters.

Step 4 Click **Submit**. The Incoming Mail Policies table page is re-displayed. Note that the summary link for the anti-spam security service has changed to reflect the new values.

Similar to the steps above, you can change the default anti-virus and virus outbreak filter settings for the default policy.

Figure 6-6 Anti-Spam Settings Page

Mail Policies: Anti-Spam

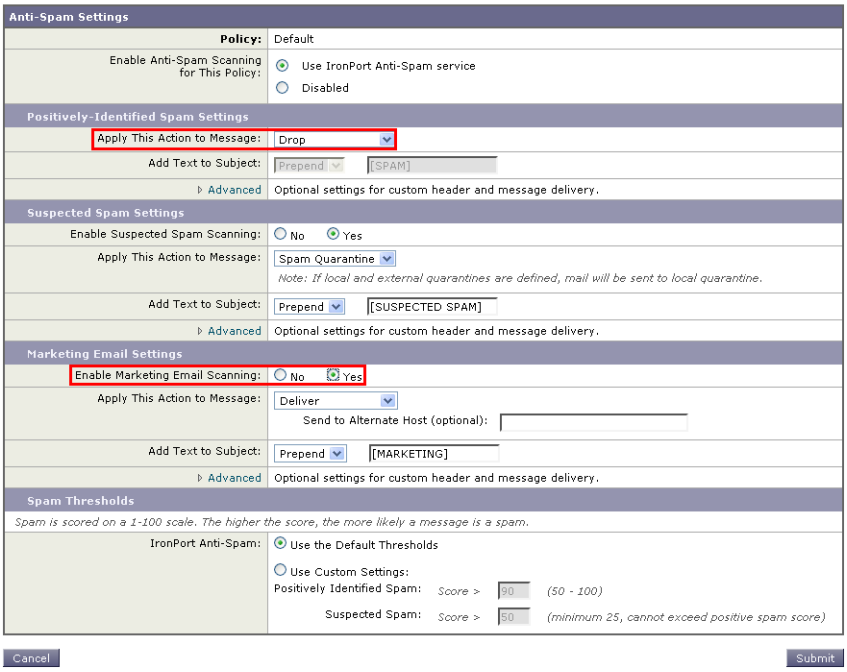

Creating a New Policy

In this part of the example, you will create two new policies: one for the sales organization (whose members will be defined by an LDAP acceptance query), and another for the engineering organization. Both policies will be assigned to the Policy Administrator custom user role to make delegated administrators belonging to this role responsible for managing these policies. You will then configure different email security settings for each.

Step 1 Click the **Add Policy** button to begin creating a new policy.

The Add Users page is displayed.

Step 2 Define a unique name for and adjust the order of the policy (if necessary).

The name of the policy must be unique to the Mail Policies table (either incoming or outgoing) in which it is defined.

Remember that each recipient is evaluated for each policy in the appropriate table (incoming or outgoing) in a top-down fashion. See First Match Wins, page 6-3 for more information.

Step 3 Click the Editable by (Roles) link and select the custom user roles for the delegated administrators who will be responsible for managing the mail policy.

When you click the link, AsyncOS displays the custom roles for delegated administrators that have edit privileges for mail policies. Delegated administrators can edit a policy's Anti-Spam, Anti-Virus, and Outbreak Filters settings and enable or disable content filters for the policy. Only operators and administrators can modify a mail policy's name or its senders, recipients, or groups. Custom user roles that have full access to mail policies are automatically assigned to mail policies.

See the "Common Administrative Tasks" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide* for more information on delegated administration.

Step 4 Define users for the policy.

You define whether the user is a sender or a recipient. (See Policy Matching, page 6-3 for more detail.) The form shown in Figure 6-7 defaults to recipients for incoming mail policies and to senders for outgoing mail policies.

Users for a given policy can be defined in the following ways:

- Full email address: `user@example.com`
- Partial email address: `user@`
- All users in a domain: `@example.com`
- All users in a partial domain: `@.example.com`
- By matching an LDAP Query

Note Entries for users are case-insensitive in both the GUI and CLI in AsyncOS. For example, if you enter the recipient `Joe@` for a user, a message sent to `joe@example.com` will match.

If you store user information within LDAP directories in your network infrastructure — for example, in Microsoft Active Directory, SunONE Directory Server (formerly known as "iPlanet Directory Server"), or Open LDAP directories — you can configure the Cisco IronPort appliance to query your LDAP servers for the purposes of accepting recipient addresses, rerouting messages to alternate addresses and/or mail hosts, masquerading headers, and determining if messages have recipients or senders from specific groups.

If you have configured the appliance to do so, you can use the configured queries to define users for a mail policy in Email Security Manager.

See the "LDAP Queries" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information.

Figure 6-7 Defining Users for a Policy

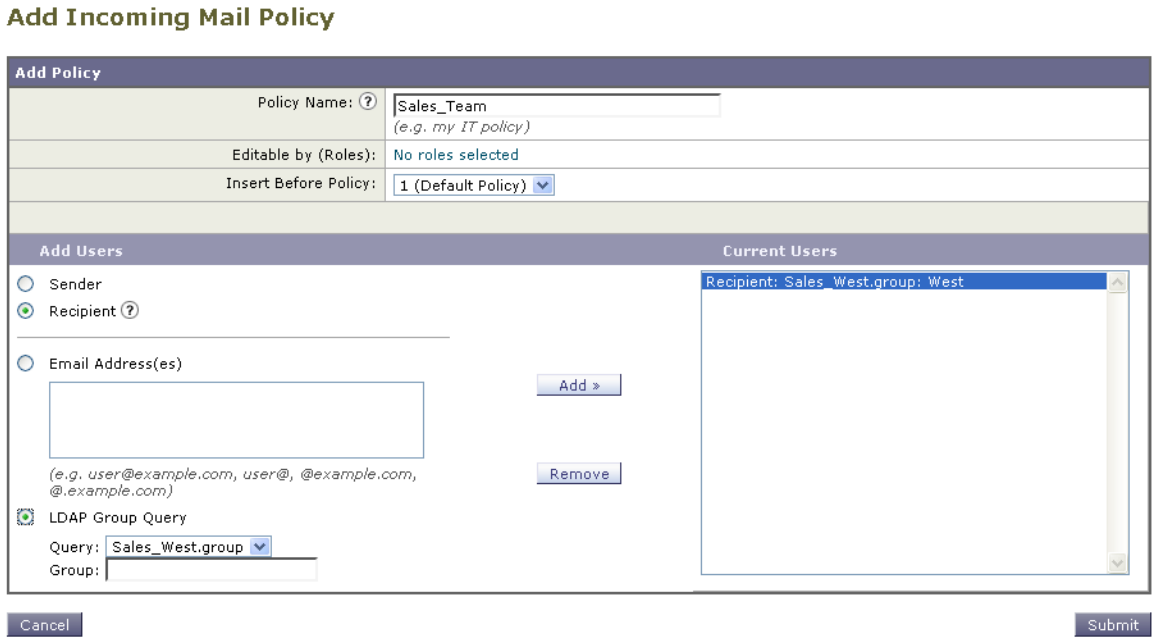

Step 5 Click the **Add** button to add users into the Current Users list.

Policies can contain mixtures of senders, recipients, and LDAP queries.

Use the **Remove** button to remove a defined user from the list of current users.

Step 6 When you are finished adding users, click **Submit**.

The Mail Policies page is displayed with the new policy added.

Note that all security services settings are set to use the default values when you first add a policy.

Figure 6-8 Newly Added Policy — Sales Group

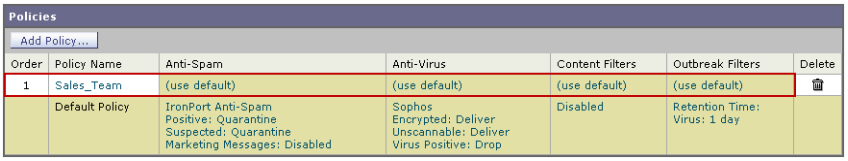

Step 7 Click the **Add Policy** button again to add another new policy.

In this policy, individual email addresses for members of the engineering team are defined:

Figure 6-9 Creating a Policy for the Engineering Team

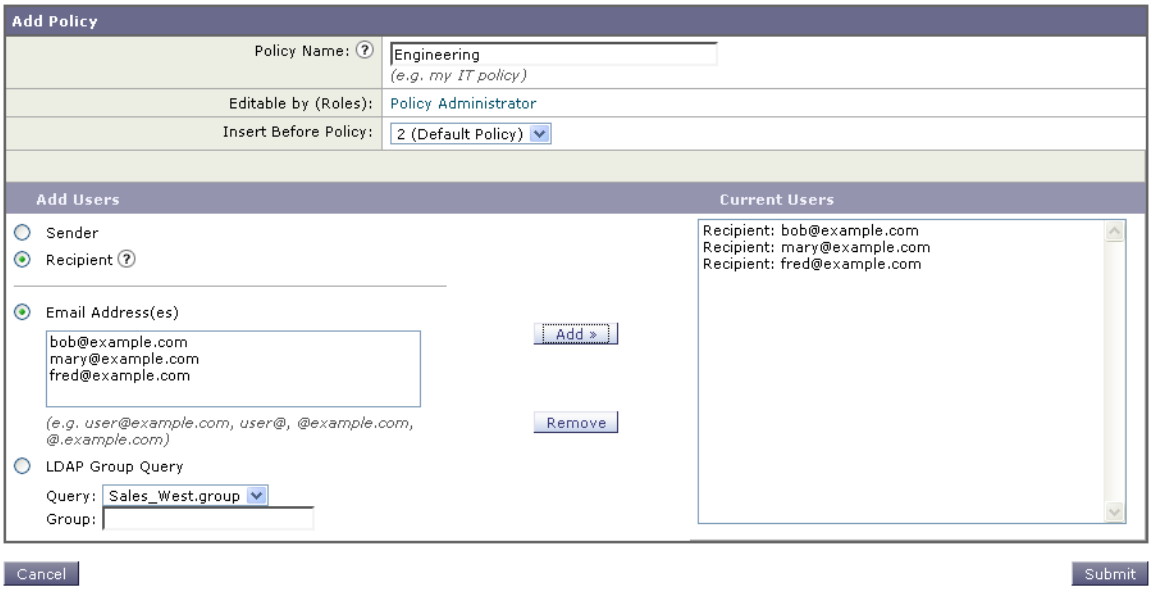

Step 8 When you are finished adding users for the engineering policy, click **Submit**.

The Mail Policies page is displayed with the new policy added. See Figure 6-10.

Step 9 Commit your changes.

Figure 6-10 Newly Added Policy — Engineering Team

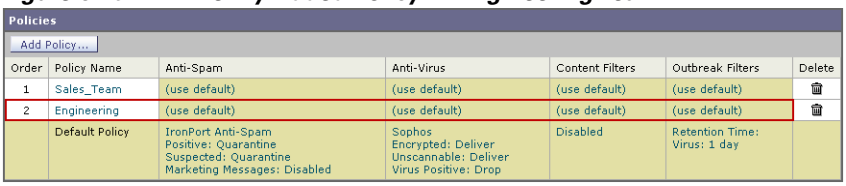

Note At this point, both newly created policies have the same settings applied to them as those in the default policy. Messages to users of either policy will match; however, the mail processing settings are not any different from the default policy. Therefore, messages that match users in the "Sales_Group" or "Engineering" policies will not be processed any differently than the default policy.

Default, Custom, and Disabled

The key at the bottom of the table shows how the color coding of cells for specific policies relates to the policy defined for the default row:

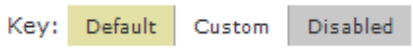

- Yellow shading shows that the policy is using the same settings as the default policy.
- No shading (white) shows that the policy is using different settings than the default policy.
- Grey shading shows that the security service has been disabled for the policy.

Creating Custom Policies

In this part of the example, you will edit the two policies just created in the previous section.

- For the sales group, you will change the anti-spam settings to be even more aggressive than the default policy. (See Editing the Default Policy: Anti-Spam Settings, page 6-21.) The default policy of dropping positively identified spam messages will be kept. However, in this example, you will change the setting for marketing messages so that they will be sent to the Cisco IronPort Spam quarantine.

  This aggressive policy has the effect of minimizing unwanted messages being sent to sales team inboxes.

  See Anti-Spam, page 9-1 for more information on anti-spam settings.

- For the engineering team, customize the Outbreak Filters feature setting so that it will modify the URLs in suspicious messages, except for links to example.com. Attachment files with the extension "dwg" will be bypassed by the Outbreak Filter scanning.

  See Outbreak Filters, page 10-1 for more information on configuring Outbreak Filters.

To edit the anti-spam settings for the sales team policy:

Step 1 Click the link for the Anti-Spam security service (the Anti-Spam) column in the sales policy row.

Because the policy was just added, the link is named: `(use default)`.

Figure 6-11 Editing the Anti-Spam Settings for the Sales Team Policy

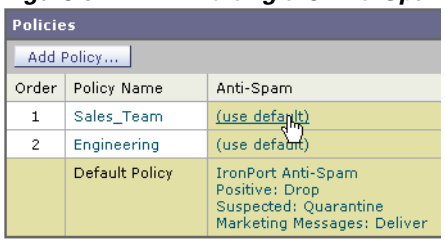

The anti-spam settings page is displayed.

Step 2 On the anti-spam security service page, change the value for "Enable Anti-Spam Scanning for this Policy" from "Use Default Settings" to "Use Cisco IronPort Anti-Spam service."

Choosing "Use Cisco IronPort Anti-Spam service" here allows you to override the settings defined in the default policy.

Step 3 In the "Positively-Identified Spam Settings" section, change the "Apply This Action to Message" to "Drop."

Step 4 In the "Suspected Spam Settings" section, click **Yes** to enable suspected spam scanning.

Step 5 In the "Suspected Spam Settings" section, change the "Apply This Action to Message" to "Spam Quarantine."

Note Selecting the Cisco IronPort Spam quarantine forwards mail according to the settings defined in the "Quarantines" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.

Step 6 In the "Add text to subject" field, click **None**.

Messages delivered to the Cisco IronPort Spam quarantine will have no additional subject tagging.

Step 7 In the "Marketing Email Settings" section, click **Yes** to enable scanning for marketing mail from legitimate sources.

Step 8 In the "Apply This Action to Message" section, select "Spam Quarantine."

Step 9 Submit and commit your changes.

The Incoming Mail Policies page is displayed with the changes shown for the sales policy. See Figure 6-12. Not that the shading shows that the policy is using different settings than the default policy.

Figure 6-12 Anti-Spam Settings for the Sales Group Policy Changed

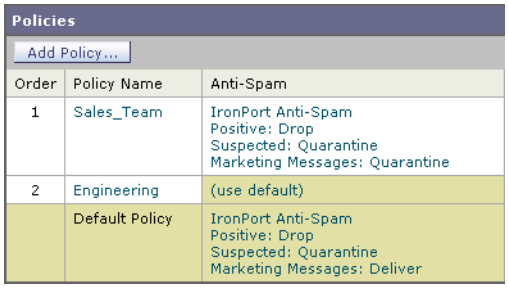

At this point, any message that is suspected spam and whose recipient matches the LDAP query defined for the sales team policy will be delivered to the Cisco IronPort Spam Quarantine.

To edit the Outbreak Filter settings for the engineering team policy:

Step 1 Click the link for the Outbreak Filters feature security service (the Outbreak Filters column) in the engineering policy row.

Because the policy was just added, the link is named: `(use default)`.

Figure 6-13 Editing the Outbreak Filters Feature Settings for the Engineering Team Policy

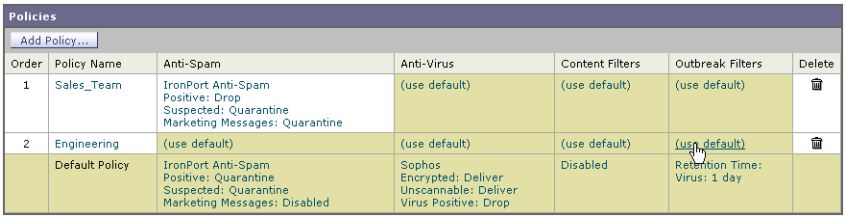

Step 2 On the Outbreak Filters feature security service page, change the scanning setting for the policy to "Enable Outbreak Filtering (Customize settings)."

Choosing "(Customize settings)" here allows you to override the settings defined in the default policy.

Doing so will also enable the contents of the rest of the page to allow you to select different settings.

Step 3 In the "Bypass Attachment Scanning" section of the page, type **`dwg`** in the in the file extension field.

The file extension "`dwg`" is not in the list of known file type that the Cisco IronPort appliance can recognize by its fingerprint when attachment scanning.

Note You do not need to type the period (`.`) before the three letter filename extension.

Step 4 Click **Add Extension** to add `.dwg` files to the list of file extensions that will bypass Outbreak Filters feature scanning.

Step 5 Click **Enable Message Modification**.

Enabling message modification allows the appliance to scan for targeted threats, such as phishing and scams, and URLs to suspicious or malicious websites. The appliance can rewrite links in messages to redirect the user through the Cisco Security proxy if they attempt to access the website.

Note Anti-spamming scanning must be enabled on the mail policy in order for Outbreak Filters to scan for targeted, non-viral threats.

Step 6 Select for **Enable for Unsigned Messages**.

This allows the appliance to rewrite URLs in signed messages. You must enable URL rewriting to be able to configure other Message Modification settings and the length of time that messages found to be non-viral threats stay in the quarantine before being released. This example uses the default retention time of 4 hours.

Step 7 Enter example.com in the **Bypass Domain Scanning** field.

The appliance will not modify links to example.com.

Step 8 Select System Generated for the **Threat Disclaimer**.

The appliance can insert a disclaimer above the message body to warn the user about the message's contents. This example uses the system generated threat disclaimer.

Figure 6-14 Outbreak Filters Settings

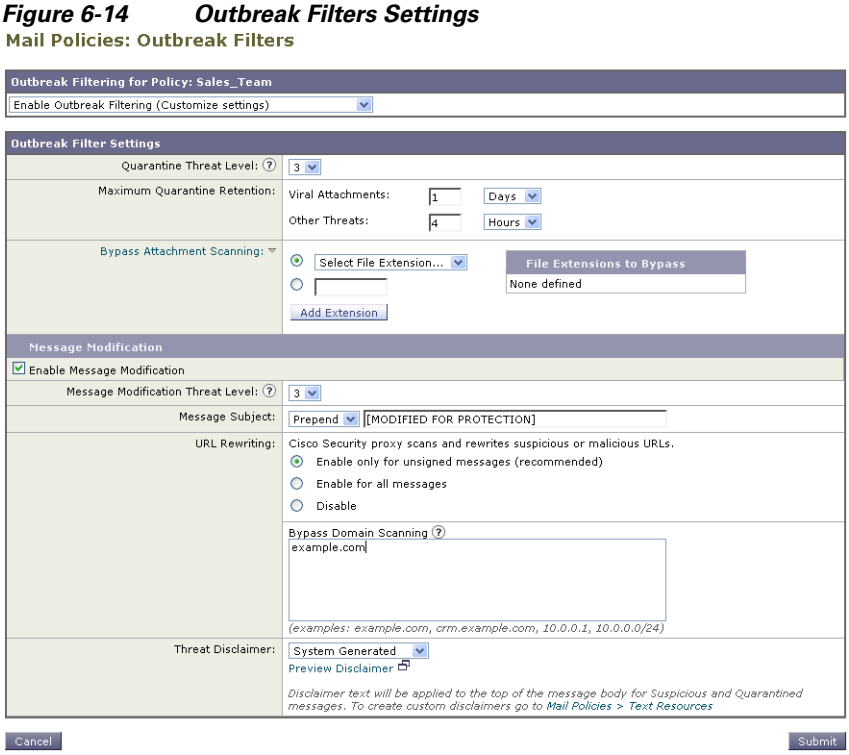

Step 9 Submit and commit your changes.

The Incoming Mail Policies page is displayed with the changes shown for the engineering policy. See Figure 6-15. Note that the shading shows that the policy is using different settings than the default policy.

Figure 6-15 Virus Filters Settings for the Engineering Policy Changed

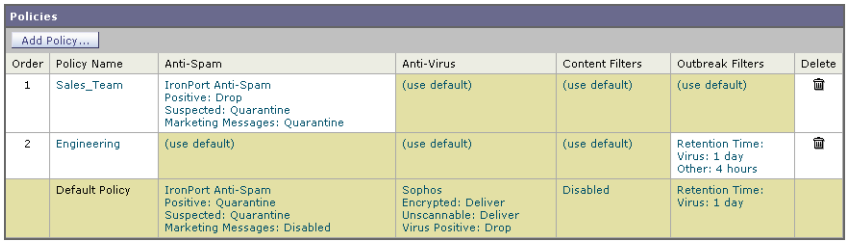

At this point, any message that contains an attachment whose file extension is `dwg` — and whose recipient matches the recipients defined for the engineering team policy — will bypass the Outbreak Filter scanning and continue processing. Messages that contain links to the example.com domain will not have their links modified to redirect through the Cisco Security proxy and will not be considered suspicious.

Finding Users in Policies of the Email Security Manager

Use the "Find Policies" button to search for users already defined in policies defined in the Email Security Manager Incoming or Outgoing Mail Policies pages.

For example, typing `joe@example.com` and clicking the Find Policies button will display results showing which policies contain defined users that will match the policy.

Figure 6-16 Finding Users in Policies

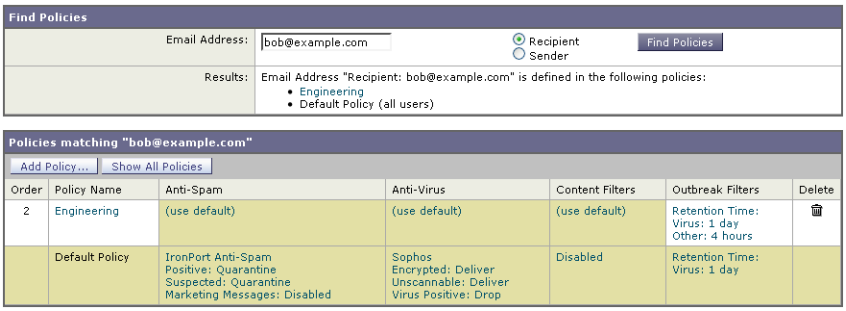

Click the name of the policy to jump to the Edit Policy page to edit the users for that policy.

Note that the default policy will always be shown when you search for any user, because, by definition, if a sender or recipient does not match any other configured policies, it will *always* match the default policy.

Email Security Manager: Managed Exceptions

Using the steps shown in the two examples above, you can begin to create and configure policies on a *managed exception* basis. In other words, after evaluating your organization's needs you can configure policies so that the majority of messages will be handled by the default policy. You can then create additional "exception" policies for specific users or user groups, managing the differing policies as needed. In this manner, message splintering will be minimized and you are less likely to impact system performance from the processing of each splinter message in the work queue.

You can define policies based on your organizations' or users' tolerance for spam, viruses, and policy enforcement. Table 6-6 on page 6-29 outlines several example policies. "Aggressive" policies are designed to minimize the amount of spam and viruses that reach end-users mailboxes. "Conservative" policies are tailored to avoid false positives and prevent users from missing messages, regardless of policies.

Table 6-6 Aggressive and Conservative Email Security Manager Settings

	Aggressive Settings	Conservative Settings
Anti-Spam	Positively identified spam: Drop Suspected spam: Quarantine Marketing mail: Deliver and prepend "`[Marketing]`" to the subject messages	Positively identified spam: Quarantine Suspected spam: Deliver and prepend "`[Suspected Spam]`" to the subject of messages Marketing mail: Disabled
Anti-Virus	Repaired messages: Deliver Encrypted messages: Drop Unscannable messages: Drop Infectious messages: Drop	Repaired messages: Deliver Encrypted messages: Quarantine Unscannable messages: Quarantine Infectious messages: Drop
Virus Filters	Enabled, no specific filename extensions or domains allowed to bypass Enable message modification for all messages	Enabled with specific filename extensions or domains allowed to bypass Enable message modification for unsigned messages

Creating New Content Filters

In this part of the example, you will create three new content filters to be used in the Incoming Mail Policy table. All of these content filters will be editable by delegated administrators belonging to the Policy Administration custom user role. You will create the following:

Step 1 "scan_for_confidential"

This filter will scan messages for the string "confidential." If the string is found, a copy of the message will be sent to email alias `hr@example.com`, and the message will be sent to the Policy quarantine area.

Step 2 "no_mp3s"

This filter will strip MP3 attachments and notify the recipients that an MP3 file was stripped.

Step 3 "ex_employee"

This content filter will scan for messages sent to a specific envelope recipient address (an ex-employee). If the message matches, a specific notification message will be sent to the sender of the message and then the message will be bounced.

After creating the content filters, you will then configure each of the policies (including the default policy) to enable the specific content filters in differing combinations.

Scan for Confidential

The first example content filter contains one condition and two actions.

Step 1 Click the Mail Policies tab.

Step 2 Click Incoming Content Filters.

The Incoming Content Filters page is displayed. On newly installed or upgraded systems, no content filters are defined by default.

Figure 6-17 Incoming Content Filters Page

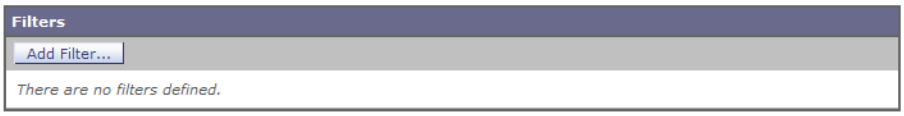

Step 3 Click the **Add Filter** button.

The Add Content Filter page is displayed.

Step 4 In the Name field, type `scan_for_confidential` as the name of the new filter.

Filter names can contain ASCII characters, numbers, underscores or dashes. The first character of a content filter name must be a letter or an underscore.

Step 5 Click the **Editable By (Roles)** link, select the Policy Administrator and click **OK**.

Delegated administrators who belong to the Policy Administrator user role will be able to edit this content filter and use it in their mail policies.

Step 6 In the Description field, type the description. For example: `scan all incoming mail for the string 'confidential'.`

Step 7 Click **Add Condition**.

Step 8 Select Message Body.

Step 9 Type `confidential` in the Contains text: field and click **OK**.

The Add Content Filter page shows the condition added.

Step 10 Click **Add Action**.

Step 11 Select Send Copy To (Bcc:).

Step 12 In the Email Addresses field, type `hr@example.com`.

Step 13 In the Subject field, type `[message matched confidential filter]`.

Step 14 Click **OK**.

The Add Content Filter page shows the action added.

Step 15 Click **Add Action**.

Step 16 Select Quarantine.

Step 17 In the drop-down menu, select the Policy quarantine area.

Step 18 Click **OK**.

The Add Content Filter page shows the second action added.

Step 19 Submit and commit your changes.

At this point, the content filter is not enabled for any incoming Mail Policy; in this example, you have only added a new content filter to the master list. Because it has not been applied to any policy, no email processing by Email Security Manager will be affected by this filter.

No MP3 Attachments

The second example content filter contains no conditions and one action.

Step 1 Click the **Add Filter** button.

The Add Content Filter page is displayed.

Step 2 In the Name field, type `no_mp3s` as the name of the new filter.

Step 3 Click the **Editable By (Roles)** link, select the Policy Administrator and click **OK**.

Step 4 In the Description field, type the description. For example: `strip all MP3 attachments`.

Step 5 Click **Add Action**.

Step 6 Select Strip Attachment by File Info.

Step 7 Select `File type is`.

Step 8 In the drop-down field, select `-- mp3`.

Step 9 Enter a replacement message if desired.

Step 10 Click **OK**.

The Add Content page shows the action added.

Step 11 Submit and commit your changes.

Note It is not necessary to specify a condition when creating a content filter. When no condition is defined, any actions defined will always apply in the rule. (Specifying no condition is equivalent to using the `true()` message filter rule — all messages will be matched if the content filter is applied to a policy.)

Ex-employee

The third content filter example uses one condition and two actions.

Step 1 Click the **Add Filter** button.

The Add Content Filter page is displayed.

Step 2 In the Name: field, type `ex_employee` as the name of the new filter.

Step 3 Click the **Editable By (Roles)** link, select the Policy Administrator and click **OK**.

Step 4 In the Description: field, type the description. For example: `bounce messages intended for Doug`.

Step 5 Click **Add Condition**.

Step 6 Select Envelope Recipient.

Step 7 For the envelope recipient, select `Begins with`, and type `doug@`.

Step 8 Click **OK**.

The Content Filters page refreshes to show the condition added. Note that you could create an LDAP directory containing the email addresses of former employees. As ex-employees are added to that directory, this content filter would be dynamically updated.

Step 9 Click **Add Action**.

Step 10 Select Notify.

Step 11 Select the checkbox for Sender and, in the Subject field, type `message bounced for ex-employee of example.com`.

Step 12 In the Use template section, select a notification template.

Note Some sections of the content filter rule builder will not appear in the user interface if the resource has not been preconfigured. For example, content dictionaries, notification templates, and message disclaimers will not appear as options if they have not been configured previously via the Mail Policies > Dictionaries page (or the `dictionaryconfig` command in the CLI). For more information about creating dictionaries, see Content Dictionaries, page 14-2.

Step 13 Click **OK**.

The Add Content Filters page shows the action added.

Step 14 Click **Add Action**.

Step 15 Select Bounce (Final Action) and click **OK**.

You can only specify one final action for a content filter. If you try to attempt to add more than one final action, the GUI displays an error.

Adding this action may will cause senders of messages to this ex-employee to potentially receive two messages: one for the notification template, and one for the bounce notification template.

Step 16 Submit and commit your changes.

The Incoming Content Filters page is displayed to show the newly-added content filter.

Enabling and Applying Content Filters to Individual Policies

In the examples above, you created three content filters using the Incoming Content Filters pages. The Incoming Content Filters and Outgoing Content filters pages hold the "master lists" of all possible content filters that can be applied to a policy.

Figure 6-18 Incoming Content Filters: Three Filters Created

In this part of the example, you will apply the three new content filters to be used in the Incoming Mail Policy table.

- The default policy will receive all three content filters.
- The engineering group will *not* receive the no_mp3s filter.
- The sales group will receive the content filters as the default incoming mail policy.

Click the links to enable and select content filters for individual policies.

Step 1 Click Incoming Mail Policies to return to the Incoming Mail Policy table.

The page is refreshed to show the default policy and the two policies added in Creating a New Policy, page 6-22. Note that content filtering is disable by default for all policies.

Step 2 Click the link for the Content Filters security service (the Content Filters column) in the default policy row. See Figure 6-19.

Figure 6-19 Editing the Content Filters Setting for the Default Incoming Mail Policy

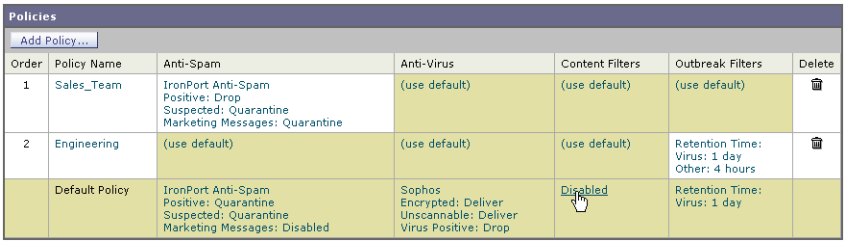

Step 3 On the Content Filtering security service page, change the value Content Filtering for Default Policy from "Disable Content Filters" to "Enable Content Filters (Customize settings)."

Figure 6-20 Enabling Content Filters for the Policy and Selecting Specific Content Filters

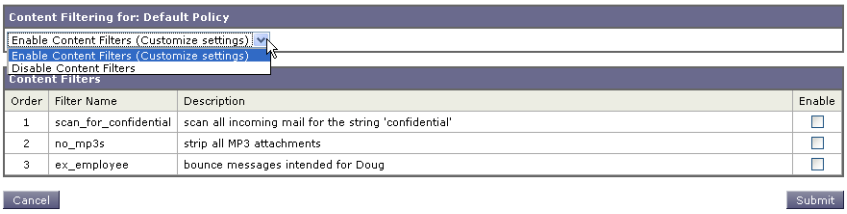

The content filters defined in the master list (which were created in Content Filters Overview, page 6-6 using the Incoming Content Filters pages) are displayed on this page. When you change the value to "Enable Content Filters (Customize settings)," the checkboxes for each filter change from disabled (greyed out) to become enabled.

Step 4 Check the **Enable** checkbox for each content filter.

Step 5 Click **Submit**.

The Incoming Mail Policies page is displayed, and the table is updated to show the names of the filters that have been enabled for the default policy.

Figure 6-21 Three Content Filters Enabled for the Default Incoming Mail Policy

Default Policy	IronPort Anti-Spam Positive: Quarantine Suspected: Quarantine Marketing Messages: Disabled	Sophos Encrypted: Deliver Unscannable: Deliver Virus Positive: Drop	scan_for_confidential no_mp3s ex_employee

To disable the "no_mp3s" content filters for the "engineering" policy:

Step 1 Click the link for the Content Filters security service (the Content Filters column) in the engineering team policy row.

Step 2 On the Content Filtering security service page, change the value for Content Filtering for Policy: Engineering from "Enable Content Filtering (Inherit default policy settings)" to "Enable Content Filtering (Customize settings)."

Because this policy was using the default values, when you change the value from "Use Default Settings" to "Yes," the checkboxes for each filter change from disabled (greyed out) to become enabled.

Step 3 Deselect the checkbox for the "no_mp3s" filter.

Figure 6-22 Deselecting a Content Filter

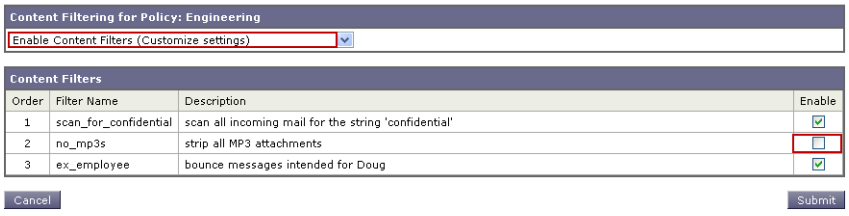

Step 4 Click **Submit**.

The Incoming Mail Policies page is displayed, and the table is updated to show the names of the filters that have been enabled for the engineering policy.

Figure 6-23 Updated Content Filters for Incoming Mail Policies

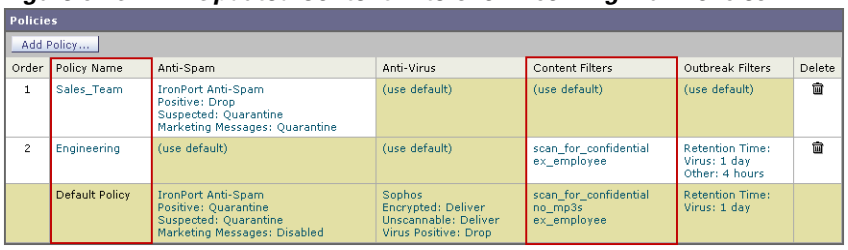

Step 5 Commit your changes.

At this point, incoming messages that match the user list for the engineering policy will not have MP3 attachments stripped; however, all other incoming messages will have MP3 attachments stripped.

Notes on Configuring Content Filters in the GUI

- It is not necessary to specify a condition when creating a content filter. When no action is defined, any actions defined will always apply in the rule. (Specifying no action is equivalent to using the `true()` message filter rule — all messages will be matched if the content filter is applied to a policy.)
- If you do not assign a custom user role to a content filter, the content filter is public and can be used by any delegated administrator for their mail policies. See the "Common Administrative Tasks" in the *Cisco IronPort AsyncOS for Email Daily Management Guide* for more information on delegated administrators and content filters.
- Administrators and operators can view and edit all content filters on an appliance, even when the content filters are assigned to custom user roles.
- When entering text for filter rules and actions, the following meta characters have special meaning in regular expression matching: `. ^ $ * + ? { [ ] \ | ( )`

  If you do not wish to use regular expression you should use a '\' (backslash) to escape any of these characters. For example: "*Warning*"
- When you define more than one Condition for a content filter, you can define whether *all* of the defined actions (that is, a logical AND) or any of the defined actions (logical OR) need to apply in order for the content filter to be considered a match.

Figure 6-24 Choosing Any or All of the Following Conditions

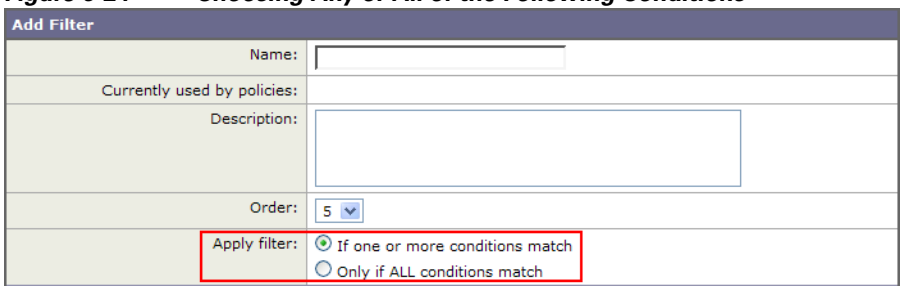

- You can test message splintering and content filters by creating "benign" content filters. For example, it is possible to create a content filter whose only action is "deliver." This content filter will not affect mail processing; however, you can use this filter to test how Email Security Manager policy processing affects other elements in the system (for example, the mail logs).
- Conversely, using the "master list" concept of the Incoming or Outgoing Content Filters, it is possible to create very powerful, wide-sweeping content filters that will immediately affect message processing for all mail handled by the appliance. The process for this is to:

  - Use the Incoming or Outgoing Content Filters page to create a new content filter whose order is 1.
  - Use the Incoming or Outgoing Mail Policies page to enable the new content filter for the default policy.
  - Enable the content filter for all remaining policies.

- The Bcc: and Quarantine actions available in Content Filters can help you determine the retention settings of quarantines you create. (See the "Quarantines" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide* for more details.) You can create filters that would simulate mail flow into and out of your system quarantines so that messages are not released too quickly from the system (that is, the quarantine areas do not fill their allotted disk space too quickly).
- Because it uses the same settings as the `scanconfig` command, the "Entire Message" condition does not scan a message's headers; choosing the "Entire Message" will scan only the message body and attachments. Use the "Subject" or "Header" conditions to search for specific header information.
- Configuring users by LDAP query will only appear in the GUI if you have LDAP servers configured on the appliance (that is, you have configured the appliance to query specific LDAP servers with specific strings using the `ldapconfig` command).
- Some sections of the content filter rule builder will not appear in the GUI if the resource has not been preconfigured. For example, notification templates and message disclaimers will not appear as options if they have not been configured previously using the Text Resources page or the `textconfig` command in the CLI.
- Content filters features will recognize, can contain, and/or scan for text in the following character encodings:
  - Unicode (UTF-8)
  - Unicode (UTF-16)
  - Western European/Latin-1 (ISO 8859-1)
  - Western European/Latin-1 (Windows CP1252)
  - Traditional Chinese (Big 5)
  - Simplified Chinese (GB 2312)
  - Simplified Chinese (HZ GB 2312)
  - Korean (ISO 2022-KR)
  - Korean (KS-C-5601/EUC-KR)
  - Japanese (Shift-JIS (X0123))
  - Japanese (ISO-2022-JP)
  - Japanese (EUC)

  You can mix and match multiple character sets within a single content filter. Refer to your web browser's documentation for help displaying and entering text in multiple character encodings. Most browsers can render multiple character sets simultaneously.

Figure 6-25 Multiple Character Sets in a Content Filter

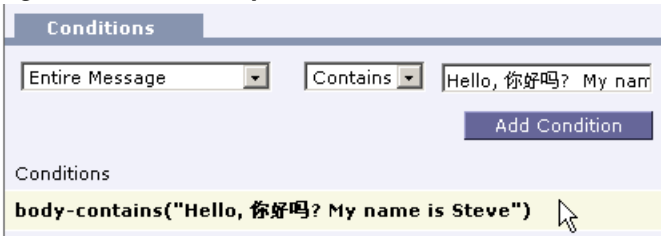

- On the Incoming or Outgoing Content Filters summary pages, use the links for "Description," "Rules," and "Policies" to change the view presented for the content filters:
  - The **Description** view shows the text you entered in the description field for each content filter. (This is the default view.)
  - The **Rules** view shows the rules and regular expressions build by the rule builder page.
  - The **Policies** shows the policies for which each content filter is enabled.

Figure 6-26 Using the Links to Toggle Description, Rules, and Policy for Content Filters

Incoming Content Filters

Filters

Add Filter...

Order	Filter Name	Description | Rules | Policies	Duplicate	Delete
1	scan_for_confidential	scan_for_confidential: if (body-contains("confidential")) { quarantine ("Policy"); bcc ("hr@example.com", "[message matched confidential filter]"); }		
2	no_mp3s	no_mp3s: if (true) { drop-attachments-by-filetype("mp3", "mp3 deleted"); }		
3	ex_employee	ex_employee: if (rcpt-to == "^doug@") { notify-copy ("$EnvelopeSender", "message bounced for ex-employee of example.com"); bounce(); }		
4	drop_large_attachments	drop_large_attachments: if (true) { drop-attachments-by-size(5242880, "This attachment was too big!"); }		

CHAPTER 7

Reputation Filtering

The Cisco IronPort appliance offers a unique, layered approach to stopping spam at the email gateway. The first layer of spam control, reputation filtering, allows you to classify email senders and restrict access to your email infrastructure based on senders' trustworthiness as determined by the Cisco IronPort SenderBase™ Reputation Service. The second layer of defense (discussed in the next chapter), scanning, is powered by Cisco IronPort Anti-Spam™ technology. Coupled together, reputation filtering and anti-spam scanning offer the most effective and highest performing anti-spam solution available today.

Using the Cisco IronPort appliance, it is very easy to create policies to deliver messages from known or highly reputable senders — such as customers and partners — directly to the end user without any anti-spam scanning. Messages from unknown or less reputable senders can be subjected to anti-spam scanning, and you can also throttle the number of messages you are willing to accept from each sender. Email senders with the worst reputation can have their connections rejected or their messages bounced based on your preferences.

The unique, two-layer approach to fighting spam of the Cisco IronPort appliance provides you with a powerful and unprecedented flexibility to manage and protect your enterprise email gateway.

Note Starting in AsyncOS 7.6, an Email Security appliance requires an anti-spam system feature key in order to use the SenderBase Reputation Service.

Reputation Filtering

The SenderBase 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 SenderBase Reputation Service returns a score based on the probability that a message from a given source is spam and exposes objective data in the Mail Flow Monitor feature to allow mail administrators to get a more complete picture of who is sending them email (see "Using Email Security Monitor" in the *Cisco IronPort AsyncOS for Email Daily Management Guide*). The SenderBase Reputation Service can is primarily designed to improve the effectiveness of a content-based anti-spam system such as Cisco IronPort Anti-Spam and requires anti-spam to be enabled on the service in order to use it.

Using the SenderBase Reputation Service, you can:

- Reduce spam

The SenderBase Reputation Service allows enterprises to identify known spam based on the connecting IP address, allowing organizations to block spam as soon as it reaches the gateway. This increases the effectiveness of the anti-spam scanning engine being used or any content-based filter.

- Protect against spam floods

Viruses such as SoBig and “hit and run” spam attacks can create sudden and unexpected spikes in message volume. If a particular sender starts sending at high volumes, the SenderBase Reputation Service can detect this through its global affiliate network and assign a more negative score, which the Cisco IronPort appliance can use to immediately begin limiting the number of recipients per hour allowed from the sender. (See also Outbreak Filters, page 10-1.)

- Improve throughput

The Cisco IronPort appliance can reduce system load and increase message throughput by immediately rejecting known spam and routing known good messages past content filters.

Reputation Filtering: the Cisco IronPort SenderBase Reputation Service

The Cisco IronPort SenderBase Reputation Service (available at `http://www.senderbase.org`) is a service designed to help email administrators better manage incoming email streams by providing objective data about the identity of senders. The SenderBase Reputation Service is similar to a credit reporting service for email; it provides data that enterprises can use to differentiate legitimate senders from spam sources. Integrated directly into the Cisco IronPort appliance GUI, the SenderBase Reputation Service provides objective data that allows you to identify reliably and block IP addresses originating unsolicited commercial email (UCE) or to verify the authenticity of legitimate incoming email from business partners, customers, or any other important source. The SenderBase 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.

Note If your Cisco IronPort appliance is set to receive mail from a local MX/MTA, you must identify upstream hosts that may mask the sender's IP address. See Incoming Relays, page 9-19 for more information.

Several key elements of the SenderBase Reputation Service are that it is:

- Non-spoofable

The email sender’s reputation is based on the IP addresses of the email sender. Because SMTP is a two-way conversation over TCP/IP, it is nearly impossible to “spoof” an IP address — the IP address presented must actually be controlled by the server sending the message.

- Comprehensive

The SenderBase Reputation Service uses global data from the SenderBase Affiliate network such as complaint rates and message volume statistics as well as data from carefully selected public blacklists and open proxy lists to determine the probability that a message from a given source is spam.

- Configurable

Unlike other “identity-based” anti-spam techniques like blacklists or whitelists that return a simple yes/no decision, the SenderBase Reputation Service returns a graduated response based on the probability that a message from that source is spam. This allows you to set your own threshold for where you choose to block spam and automatically assign senders to different groups based on their SenderBase Reputation Score.

SenderBase Reputation Score (SBRS)

The SenderBase Reputation Score (SBRS) is a numeric value assigned to an IP address based on information from the SenderBase Reputation Service. The SenderBase Reputation Service aggregates data from over 25 public blacklists and open proxy lists, and combines this data with global data from SenderBase to assign a score from -10.0 to +10.0, as follows:

Score	Meaning
-10.0	Most likely to be a source of spam
0	Neutral, or not enough information to make a recommendation
+10.0	Most likely to be a trustworthy sender

The lower (more negative) the score, the more likely that a message is spam. A score of -10.0 means that this message is "guaranteed" to be spam, while a score of 10.0 means that the message is "guaranteed" to be legitimate.

Using the SBRS, you configure the Cisco IronPort appliance to apply mail flow policies to senders based on their trustworthiness. (You can also create message filters to specify "thresholds" for SenderBase Reputation Scores to further act upon messages processed by the system. For more information, refer to "SenderBase Reputation Rule" and "Bypass Anti-Spam System Action" in the "Using Message Filters to Enforce Email Policies" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.)

Figure 7-1 The SenderBase Reputation Service

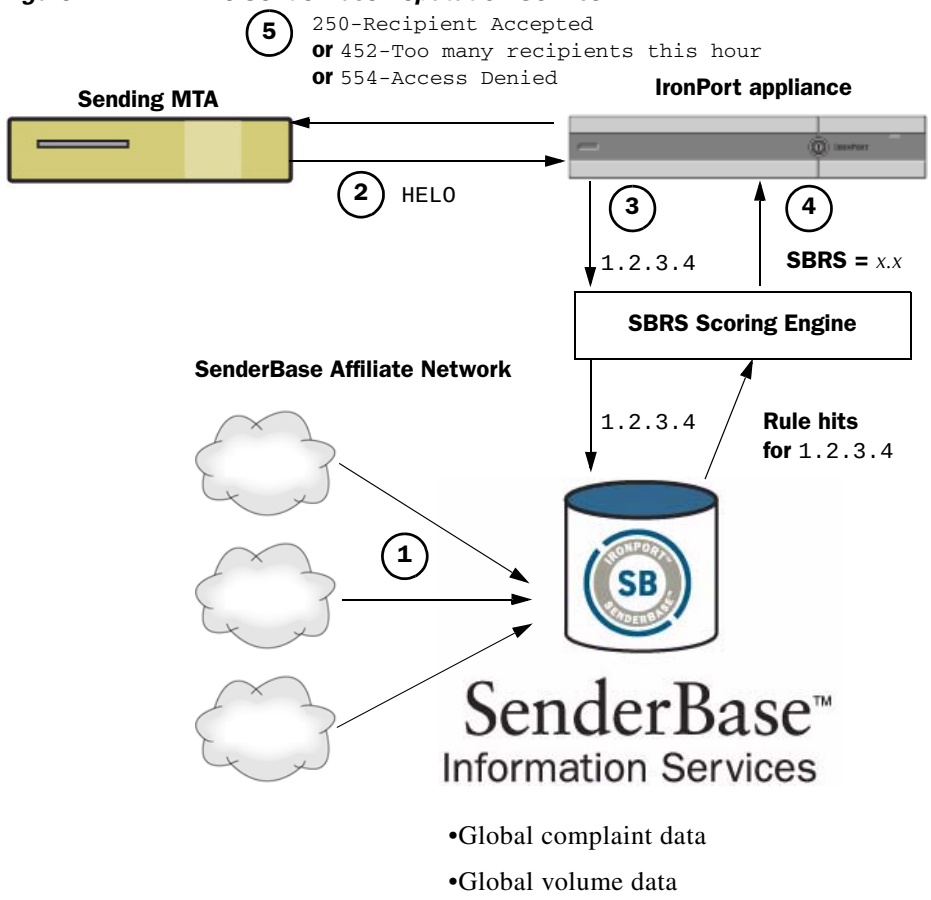

Step 1 SenderBase affiliates send real-time, global data

Step 2 Sending MTA opens connection with the Cisco IronPort appliance

Step 3 Cisco IronPort appliance checks global data for the connecting IP address

Step 4 SenderBase Reputation Service calculates the probability this message is spam and assigns a SenderBase Reputations Score

Step 5 Cisco IronPort returns the response based on the SenderBase Reputation Score

Implementing SenderBase Reputation Filters

Cisco IronPort Reputation Filter technology aims to shunt as much mail as possible from the remaining security services processing that is available on the Cisco IronPort appliance. (See Understanding the Email Pipeline, page 4-1.)

When enabling reputation filtering, mail from known bad senders is simply refused. Known good mail from global 2000 companies is automatically routed around the spam filters, reducing the chance of false positives. Unknown, or “grey” email is routed to the anti-spam scanning engine. Using this approach, reputation filters can reduce the load on the content filters by as much as 50%.

Figure 7-2 Reputation Filtering Example

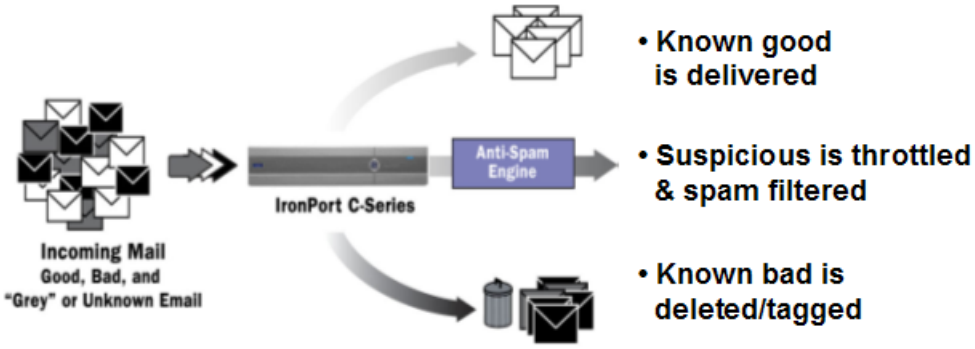

Table 7-2 lists a set of recommended policies for implementing SenderBase reputation filtering. Depending on the objectives of your enterprise, you can implement a conservative, moderate, or aggressive approach.

Note Although Cisco recommends throttling, an alternative for implementing the SenderBase Reputation Service is to modify the subject line of suspected spam messages. To do this, use the following message filter shown in Table 7-1. This filter uses the `reputation` filter rule and the `strip-header` and `insert-header` filter actions to replace the subject line of messages with a SenderBase Reputation Score lower than -2.0 with a subject line that includes the actual SenderBase Reputation Score represented as: **`{Spam`** *SBRS*`}`. Replace *listener_name* in this example with the name of your public listener. (The period on its own line is included so that you can cut and paste this text directly into the command line interface of the `filters` command.)

Refer to "Using Message Filters to Enforce Email Policies" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*. for more information.

Table 7-1 Message Filter to Modify Subject Header with SBRS: Example 1

```
sbrs_filter:

if ((recv-inj == "listener_name" AND subject != "\\{Spam -?[0-9.]+\\}"))

{

        insert-header("X-SBRS", "$REPUTATION");

        if (reputation <= -2.0)

{

        strip-header("Subject");

        insert-header("Subject", "$Subject \\{Spam $REPUTATION\\}");

      }

 }

.
```

Configuring Reputation Filtering

Configure reputation filtering via the Mail Policies > HAT Overview page. For more information, see Implementing SenderBase Reputation Filters, page 7-4.

Conservative

A conservative approach is to block messages with a SenderBase Reputation Score lower than -4.0, throttle between -4.0 and -2.0, apply the default policy between -2.0 and +6.0, and apply the trusted policy for messages with a score greater than +6.0. Using this approach ensures a near zero false positive rate while achieving better system performance.

Moderate

A moderate approach is to block messages with a SenderBase Reputation Score lower than -3.0, throttle between -3.0 and 0, apply the default policy between 0 and +6.0, and apply the trusted policy for messages with a score greater than +6.0. Using this approach ensures a very small false positive rate while achieving better system performance (because more mail is shunted away from Anti-Spam processing).

Aggressive

An aggressive approach is to block messages with a SenderBase Reputation Score lower than -2.0, throttle between -2.0 and 0.5, apply the default policy between 0 and +4.0, and apply the trusted policy for messages with a score greater than +4.0. Using this approach, you might incur some false positives; however, this approach maximizes system performance by shunting the most mail away from Anti-Spam processing.

Note Users are also recommended to assign all messages with a SenderBase Reputation Score greater than 6.0 to the $TRUSTED policy.

Table 7-2 Recommended Phased Approach to Implementing Reputation Filtering using the SBRS

Policy	Blacklist	Throttle	Default	Whitelist
Conservative	-10 to -4	-4 to -2	-2 to 7	7 to 10
Moderate	-10 to -3	-3 to -1	-1 to 6	6 to 10
Aggressive	-10 to -2	-2 to -0.5	-0.5 to 4	4 to 10

Policy:	Characteristics:	Mail Flow Policy to Apply:
Conservative:	Near zero false positives, better performance	**$BLOCKED**
Moderate:	Very few false positives, high performance	**$THROTTLED**
Aggressive:	Some false positives, maximum performance	**$DEFAULT**

The steps below outline a phased approach to implementing reputation filtering:

Implementing Reputation Filtering in a Listener's HAT

Step 1 From the Mail Policies tab, select Host Access Table > HAT Overview. Select the public listener from the Sender Groups (Listener) menu. The HAT Overview page shows the SenderBase Reputation Score settings for each Sender Group.

Figure 7-3 Listing Sender Groups' SenderBase Reputation Score Ranges

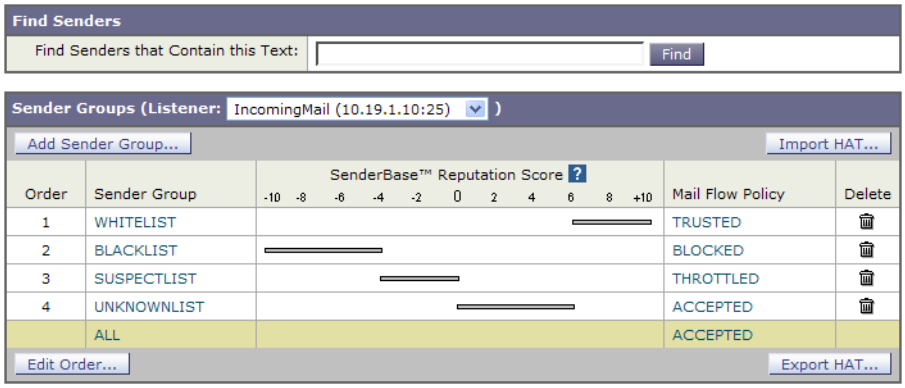

The HAT Overview shows the range of SenderBase Reputation Scores that are assigned to each sender group (the horizontal bar) as well as the associated mail flow policy.

Step 2 Click the link for a sender group.

For example, click the "SUSPECTLIST" link. The Edit Sender Group page is displayed:

Figure 7-4 Modifying a Sender Group's SBRS Ranges

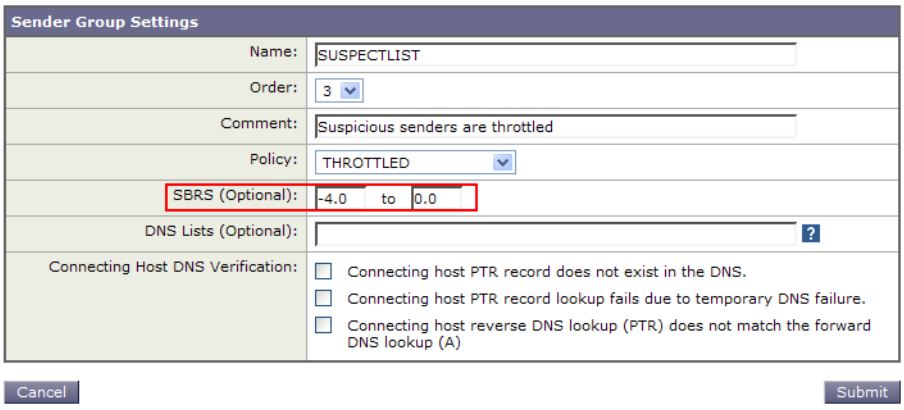

Step 3 Type the range of SenderBase Reputation Scores to define the sender group. You can also define an optional comment.

For example, for "SUSPECTLIST," enter a range from -4.0 to 0. Refer to Sender Groups defined by SenderBase Reputation Scores, page 5-23 for the syntax.

Step 4 Click **Submit**.

Repeat steps 2-5 for each group in the listener's HAT. For example, define the values for *conservative* approach. You can configure the values shown in Table 7-2 for a moderate or aggressive approach as well.

Sender Group	SBRS Range	Mail Flow Policy
WHITELIST	6 to 10	TRUSTED
BLACKLIST	-10 to -7	BLOCKED
SUSPECTLIST	-7 to -2	THROTTLED
UNKOWNLIST	-2 to 6	ACCEPTED

Note Remember that order matters when defining sender groups in a listener's HAT. (The HAT is read from top to bottom for each host that attempts to connect to the listener. If a rule matches a connecting host, the action is taken for that connection immediately.) Cisco recommends maintaining the default order of the predefined sender groups in a listener's HAT — that is, RELAYLIST (C10/100 customers only), followed by WHITELIST, BLACKLIST, SUSPECTLIST, and UNKNOWNLIST.

Step 5 Click the **Commit Changes** button, add an optional comment if necessary, and then click **Commit Changes** to finish implementing reputation filtering in a listener's HAT.

Testing Reputation Filtering Using the SBRS

Unless you regularly receive a large portion of spam, or you have set up "dummy" accounts to specifically receive spam for your organization, it may be difficult to immediately test the SBRS policies you have implemented. However, if you add entries for reputation filtering with SenderBase Reputation Scores into a listener's HAT as indicated in Table 7-3, you will notice that a smaller percentage of inbound mail will be "unclassified."

You test the policies you have created using the `trace` command with an arbitrary SBRS. See Debugging Mail Flow Using Test Messages: Trace, page -446. The `trace` command is available in the CLI as well as the GUI.

Table 7-3 Suggested Mail Flow Policies for Implementing the SBRS

Policy Name	Primary Behavior (Access Rule)	Parameters	Value
$BLOCKED	REJECT	None	
$THROTTLED	ACCEPT	Maximum messages / session:	10
		Maximum recipients / message:	20
		Maximum message size:	1 MB
		Maximum concurrent connections:	10
		Use Spam Detection:	ON
		Use TLS:	OFF
		Maximum recipients / hour:	**20** *(recommended)*
		Use SenderBase:	ON
$ACCEPTED (Public Listener)	ACCEPT	Maximum messages / session:	1,000
		Maximum recipients / message:	1,000
		Maximum message size:	100 MB
		Maximum concurrent connections:	1,000
		Use Spam Detection:	ON
		Use TLS:	OFF
		Use SenderBase:	ON
$TRUSTED	ACCEPT	Maximum messages / session:	1,000
		Maximum recipients / message:	1,000
		Maximum message size:	100 MB
		Maximum concurrent connections:	1,000
		Use Spam Detection:	OFF
		Use TLS:	OFF
		Maximum recipients / hour:	-1 *(disabled)*
		Use SenderBase:	OFF

Note In the $THROTTLED policy, the maximum recipients per hour from the remote host is set to 20 recipient per hour, by default. Note that this setting controls the maximum throttling available. You can increase the number of recipients to receive per hour if this parameter is too aggressive. For more information on Default Host Access policies, see Accessing Predefined Mail Flow Policies for Public Listeners, page 5-25.

Monitoring the Status of the SenderBase Reputation Service

The SenderBase page in the Security Services menu displays the connection status and the timestamp of the most recent query from the Cisco IronPort appliance to the SenderBase Network Status Server and SenderBase Reputation Score Service. The SenderBase Reputation Score Service sends the SRBS scores to the appliance. The SenderBase Network Server sends the appliance information the IP addresses, domains, and organizations that are sending mail to you. AsyncOS uses this data for its reporting and email monitoring features.

Figure 7-5 SenderBase Network Status on the SenderBase Page

SenderBase Network Status		
Type	Status	Last Status Check
SenderBase Network Server	up	Wed Sep 10 13:44:52 2008 PDT
SenderBase Reputation Score Service	up	Wed Sep 10 13:44:52 2008 PDT

The `sbstatus` command in CLI displays the same information.

CHAPTER 8

Anti-Virus

The Cisco IronPort appliance includes integrated virus scanning engines from Sophos, Plc and McAfee, Inc. You can obtain license keys for the Cisco IronPort appliance to scan messages for viruses using one or both of these virus scanning engines.

You can configure the appliance to scan messages for viruses (based on the matching incoming or outgoing mail policy), and, if a virus is found, to perform different actions on the message (including "repairing" the message of viruses, modifying the subject header, adding an additional X-header, sending the message to an alternate address or mailhost, archiving the message, or deleting the message).

If enabled, virus scanning is performed in the "work queue" on the appliance, immediately after Anti-Spam scanning. (See Understanding the Email Pipeline, page 4-1.)

By default, virus scanning is enabled for the default incoming and outgoing mail policies.

Anti-Virus Scanning

You can configure your Cisco IronPort appliance to scan for viruses using the McAfee or Sophos anti-virus scanning engines.

The McAfee and Sophos engines contain the program logic necessary to scan files at particular points, process and pattern-match virus definitions with data they find in your files, decrypt and run virus code in an emulated environment, apply heuristic techniques to recognize new viruses, and remove infectious code from legitimate files.

Evaluation Key

Your Cisco IronPort appliance ships with a 30-day evaluation key for each available anti-virus scanning engine. You enable the evaluation key by accessing the license agreement in the System Setup Wizard or Security Services > Sophos/McAfee Anti-Virus pages (in the GUI) or running the `antivirusconfig` or `systemsetup` commands (in the CLI). Once you have accepted the agreement, the Anti-Virus scanning

engine will be enabled, by default, for the default incoming and outgoing mail policies. For information on enabling the feature beyond the 30-day evaluation period, contact your Cisco IronPort sales representative. You can see how much time remains on the evaluation via the System Administration > Feature Keys page or by issuing the `featurekey` command. (For more information, see the section on working with feature keys in "Common Administrative Tasks" in the *Cisco IronPort AsyncOS for Email Daily Management Guide*).

Multi-Layer Anti-Virus Scanning

AsyncOS supports scanning messages with multiple anti-virus scanning engines — multi-layer anti-virus scanning. You can configure your Cisco IronPort appliance to use one or both of the licensed anti-virus scanning engines on a per mail policy basis. You could create a mail policy for executives, for example, and configure that policy to scan mail with both Sophos and McAfee engines.

Scanning messages with multiple scanning engines provides "defense in depth" by combining the benefits of both Sophos and McAfee anti-virus scanning engines. Each engine has leading anti-virus capture rates, but because each engine relies on a separate base of technology (discussed in McAfee Anti-Virus Filtering, page 8-4 and Sophos Anti-Virus Filtering, page 8-2) for detecting viruses, the multi-scan approach can be even more effective. Using multiple scanning engines can lead to reduced system throughput, please contact your Cisco IronPort support representative for more information.

You cannot configure the order of virus scanning. When you enable multi-layer anti-virus scanning, the McAfee engine scans for viruses first, and the Sophos engine scans for viruses second. If the McAfee engine determines that a message is virus-free, the Sophos engine scans the message, adding a second layer of protection. If the McAfee engine determines that a message contains a virus, the Cisco IronPort appliance skips Sophos scanning and performs actions on the virus message based on settings you configured.

Sophos Anti-Virus Filtering

The Cisco IronPort appliance includes integrated virus-scanning technology from Sophos, Plc. Sophos Anti-Virus provides cross-platform anti-virus protection, detection and disinfection.

Sophos Anti-Virus provides a virus detection engine that scans files for viruses, Trojan horses, and worms. These programs come under the generic term of *malware*, meaning "malicious software." The similarities between all types of malware allow anti-virus scanners to detect and remove not only viruses, but also all types of malicious software.

Virus Detection Engine

The Sophos virus detection engine lies at the heart of the Sophos Anti-Virus technology. It uses a proprietary architecture similar to Microsoft's COM (Component Object Model), consisting of a number of objects with well-defined interfaces. The modular filing system used by the engine is based on separate, self-contained dynamic libraries each handling a different "storage class," for example, file type. This approach allows virus scanning operations to be applied on generic data sources, irrespective of type.

Specialized technology for loading and searching data enables the engine to achieve very fast scanning speeds. Incorporated within it are:

- a full code emulator for detecting polymorphic viruses

- an on-line decompressor for scanning inside archive files
- an OLE2 engine for detecting and disinfecting macro viruses

The Cisco IronPort appliance integrates with the virus engine using SAV Interface.

Virus Scanning

In broad terms, the engine's scanning capability is managed by a powerful combination of two important components: a classifier that knows where to look, and the virus database that knows what to look for. The engine classifies the file by type rather than by relying on the extension.

The virus engine looks for viruses in the bodies and attachments of messages received by the system; an attachment's file type helps determine its scanning. For example, if a message's attached file is an executable, the engine examines the header which tells it where the executable code starts and it looks there. If the file is a Word document, the engine looks in the macro streams. If it is a MIME file, the format used for mail messaging, it looks in the place where the attachment is stored.

Detection Methods

How viruses are detected depends on their type. During the scanning process, the engine analyzes each file, identifies the type, and then applies the relevant technique(s). Underlying all methods is the basic concept of looking for certain types of instructions or certain ordering of instructions.

Pattern matching

In the technique of pattern matching, the engine knows the particular sequence of code and is looking for an exact match that will identify the code as a virus. More often, the engine is looking for sequences of code that are similar, but not necessarily identical, to the known sequences of virus code. In creating the descriptions against which files are compared during scanning, Sophos virus researchers endeavor to keep the identifying code as general as possible so that – using heuristics, as explained below – the engine will find not just the original virus but also its later derivatives.

Heuristics

The virus engine can combine basic pattern matching techniques with heuristics – a technique using general rather than specific rules – to detect several viruses in the same family, even though Sophos researchers might have analyzed only one virus in that family. The technique enables a single description to be created that will catch several variants of one virus. Sophos tempers its heuristics with other methods, minimizing the incidence of false positives.

Emulation

Emulation is a technique applied by the virus engine to polymorphic viruses. Polymorphic viruses are encrypted viruses that modify themselves in an effort to hide themselves. There is no visible constant virus code and the virus encrypts itself differently each time it spreads. When it runs, it decrypts itself. The emulator in the virus detection engine is used on DOS and Windows executables, while polymorphic macro viruses are found by detection code written in Sophos's Virus Description Language.

The output of this decryption is the real virus code and it is this output that is detected by the Sophos virus detection engine after running in the emulator.

Executables that are sent to the engine for scanning are run inside the emulator, which tracks the decryption of the virus body as it is written to memory. Normally the virus entry point sits at the front end of a file and is the first thing to run. In most cases, only a small amount of the virus body has to be decrypted in order for the virus to be recognized. Most clean executables stop emulating after only a few instructions, which reduces overhead.

Because the emulator runs in a restricted area, if the code does turn out to be a virus, the virus does not infect the appliance.

Virus Descriptions

Sophos exchanges viruses with other trusted anti-virus companies every month. In addition, every month customers send thousands of suspect files directly to Sophos, about 30% of which turn out to be viruses. Each sample undergoes rigorous analysis in the highly secure virus labs to determine whether or not it is a virus. For each newly discovered virus, or group of viruses, Sophos creates a description.

Sophos Alerts

Cisco encourages customers who enable Sophos Anti-Virus scanning to subscribe to Sophos alerts on the Sophos site at http://www.sophos.com/virusinfo/notifications/.
Subscribing to receive alerts directly from Sophos will ensure you are apprised of the latest virus outbreaks and their available solutions.

When a Virus is Found

When a virus has been detected, Sophos Anti-Virus can repair (disinfect) the file. Sophos Anti-Virus can usually repair any file in which a virus has been found, after which the file can be used without risk. The precise action taken depends on the virus.

There can be limitations when it comes to disinfecting, because it is not always possible to return a file to its original state. Some viruses overwrite part of the executable program which cannot be reinstated. In this instance, you define how to handle messages with attachments that could not be repaired. You configure these settings on a per-recipient basis using the Email Security Feature: the Mail Policies > Incoming or Outgoing Mail Policies pages (GUI) or the `policyconfig -> antivirus` command (CLI). For more information on configuring these settings, see Configuring Virus Scanning Actions for Users, page 8-8.

McAfee Anti-Virus Filtering

The McAfee® scanning engine:

- Scans files by pattern-matching virus signatures with data from your files.
- Decrypts and runs virus code in an emulated environment.
- Applies heuristic techniques to recognize new viruses.
- Removes infectious code from files.

Pattern-Matching Virus Signatures

McAfee uses anti-virus definition (DAT) files with the scanning engine to detect particular viruses, types of viruses, or other potentially unwanted software. Together, they can detect a simple virus by starting from a known place in a file, then searching for a virus signature. Often, they must search only a small part of a file to determine that the file is free from viruses.

Encrypted Polymorphic Virus Detection

Complex viruses avoid detection with signature scanning by using two popular techniques:

- **Encryption**. The data inside the virus is encrypted so that anti-virus scanners cannot see the messages or computer code of the virus. When the virus is activated, it converts itself into a working version, then executes.
- **Polymorphism**. This process is similar to encryption, except that when the virus replicates itself, it changes its appearance.

To counteract such viruses, the engine uses a technique called emulation. If the engine suspects that a file contains such a virus, the engine creates an artificial environment in which the virus can run harmlessly until it has decoded itself and its true form becomes visible. The engine can then identify the virus by scanning for a virus signature, as usual.

Heuristics Analysis

Using only virus signatures, the engine cannot detect a new virus because its signature is not yet known. Therefore the engine can use an additional technique — heuristic analysis.

Programs, documents or email messages that carry a virus often have distinctive features. They might attempt unprompted modification of files, invoke mail clients, or use other means to replicate themselves. The engine analyzes the program code to detect these kinds of computer instructions. The engine also searches for legitimate non-virus-like behavior, such as prompting the user before taking action, and thereby avoids raising false alarms.

By using these techniques, the engine can detect many new viruses.

When a Virus is Found

When a virus has been detected, McAfee can repair (disinfect) the file. McAfee can usually repair any file in which a virus has been found, after which the file can be used without risk. The precise action taken depends on the virus.

Occasionally, there can be limitations when it comes to disinfecting files because it is not always possible to return a file to its original state. Some viruses overwrite part of the executable program which cannot be reinstated. In this instance, you define how to handle messages with attachments that could not be repaired. You configure these settings on a per-recipient basis using the Email Security Feature: the Mail Policies > Incoming or Outgoing Mail Policies pages (GUI) or the `policyconfig -> antivirus` command (CLI). For more information on configuring these settings, see Configuring Virus Scanning Actions for Users, page 8-8.

Enabling Virus Scanning and Configuring Global Settings

To perform virus scanning, you must first enable virus scanning on the Cisco IronPort appliance. After you enable the virus scanning engine or engines, you can apply the virus scanning engine to incoming or outgoing mail policies.

Overview

You can enable a virus scanning engine when you run the System Setup Wizard. Or, you can enable and modify the virus scanning engine global configuration settings via Security Services > Sophos/McAfee Anti-Virus pages (GUI) or the `antivirusconfig` command (CLI). You can configure the following global settings:

- Globally enable McAfee or Sophos anti-virus scanning for the entire system.
- Specify the anti-virus scanning timeout value.

In addition to the two values on the global settings page, you can further configure the anti-virus settings via the Service Updates page (available from the Security Services tab). Additional settings include:

- How (from which URL) the system will receive anti-virus updates. The virus definitions are updated from a dynamic URL. If you have strict firewall policies, you may need to configure your Cisco IronPort appliance to obtain updates from a static URL.
- How frequently the system checks for new virus definitions. (You define the number of minutes between checks.)
- You can optionally enable a proxy server for obtaining anti-virus updates.

For more information about configuring these additional settings, see Service Updates, page 15-10.

Enabling Anti-Virus Scanning and Configure Global Settings

Step 1 Select Security Services > McAfee.

Or

Select Security Services > Sophos.

Step 2 Click **Enable**. The license agreement page is displayed.

Note Clicking **Enable** enables the feature globally for the appliance. However, you must later enable per-recipient settings in Mail Policies.

Step 3 After reading the agreement, scroll to the bottom of the page and click **Accept** to accept the agreement.

Step 4 Click **Edit Global Settings**.

Step 5 Choose a maximum virus scanning timeout value.

Configure a timeout value for the system to stop performing anti-virus scanning on a message. The default value is 60 seconds.

Step 6 Submit and commit your changes.

Step 7 You can now configure anti-virus settings on a per-recipient basis. See Configuring Virus Scanning Actions for Users, page 8-8.

Note For information about how and when anti-virus scanning is applied, see Email Pipeline and Security Services, page 4-6.

Retrieving Anti-Virus Updates via HTTP

By default, the Cisco IronPort appliance is configured to check for updates every 5 minutes. For the Sophos and McAfee anti-virus engines, the server updates from a dynamic website.

The system does not timeout on updates as long as the update is actively downloading to the appliance. If the update download pauses for too long, then the download times out.

The maximum amount of time that the system waits for an update to complete before timing out is a dynamic value that is defined as 1 minute less than the anti-virus update interval (defined on Security Services > Service Updates). This configuration value aids appliances on slower connections while downloading large updates that may take longer than 10 minutes to complete.

Monitoring and Manually Checking for Updates

Once you have accepted the license agreement and configured the global settings, you can use the Security Services > Sophos or McAfee Anti-Virus page (GUI) or the `antivirusstatus` command (CLI) to verify that you have the latest anti-virus engine and identity files installed, and to confirm when the last update was performed.

You can also manually perform updates. From the Security Services > Sophos or McAfee Anti-Virus page, click Update Now in the Current McAfee/Sophos Anti-Virus Files table. The appliance checks for and downloads the latest updates.

Figure 8-1 Manually Checking for Sophos Updates

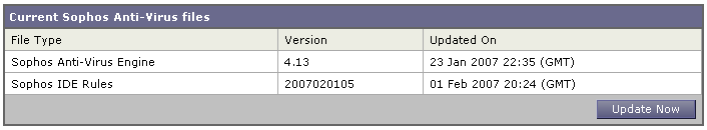

In the CLI, use the `antivirusstatus` command to check the status of your virus files and `antivirusupdate` command to manually check for updates:

```
example.com> antivirusstatus

Choose the operation you want to perform:

- MCAFEE - Display McAfee Anti-Virus version information

- SOPHOS - Display Sophos Anti-Virus version information

> sophos

SAV Engine Version        3.2.07.286_4.58
    IDE Serial                0
    Last Engine Update        Base Version
    Last IDE Update           Never updated
```

```
example.com> antivirusupdate

Choose the operation you want to perform:

- MCAFEE - Request updates for McAfee Anti-Virus

- SOPHOS - Request updates for Sophos Anti-Virus

>sophos

Requesting check for new Sophos Anti-Virus updates

example.com>
```

You can view the Updater Logs to verify whether or not the antivirus files have been successfully downloaded, extracted, or updated. Use the `tail` command to show the final entries in the Updater log subscription to ensure that virus updates were obtained.

Configuring Virus Scanning Actions for Users

Once enabled globally, the virus scanning engine integrated into the Cisco IronPort appliance processes messages for viruses for incoming and outgoing mail based on policies (configuration options) you configure using the Email Security Manager feature. You enable Anti-Virus actions on a per-recipient basis using the Email Security Feature: the Mail Policies > Incoming or Outgoing Mail Policies pages (GUI) or the `policyconfig > antivirus` command (CLI).

Message Scanning Settings

- Scan for Viruses Only:

  Messages processed by the system are scanned for viruses. Repairs are *not* attempted for infected attachments. You can choose whether to drop attachments and deliver mail for messages that contain viruses or could not be repaired.

- Scan and Repair Viruses:

  Messages processed by the system are scanned for viruses. If a virus is found in an attachment, the system will attempt to "repair" the attachment.

- Dropping Attachments

  You can choose to drop infected attachments.

  When infected attachments to messages have been scanned and *dropped* by the anti-virus scanning engine, the attachment is replaced with a new attachment called "Removed Attachment." The attachment type is text/plain and contains the following:

```
This attachment contained a virus and was stripped.

     Filename: filename

 Content-Type: application/filetype
```

Users will always be notified if their messages were modified in any way because they were infected with a bad attachment. You can configure a secondary notification action, as well (see Sending Notifications, page 8-12). The notify action is *not* needed to inform users that a message was modified if you choose to drop infected attachments.

- X-IronPort-AV Header

  All messages that are processed by the Anti-Virus scanning engine on the appliance have the header `X-IronPort-AV:` added to messages. This header provides additional information to you when debugging issues with your anti-virus configuration, particularly with messages that are considered "unscannable." You can toggle whether the X-IronPort-AV header is included in messages that are scanned. Including this header is recommended.

Message Handling Settings

You configure the virus scanning engine to handle four distinct classes of messages that are received by a listener, with separate actions for each. Figure 8-2 summarizes the actions the system performs when the virus scanning engine is enabled. See also Figure 8-3 and Figure 8-4 for the GUI configuration.

For each of the following message types, you can choose which actions are performed. The actions are described below (see Configuring Settings for Message Handling Actions, page 8-10). For example, you can configure your anti- virus settings for virus-infected messages so that the infected attachment is dropped, the subject of the email is modified, and a custom alert is sent to the message recipient.

Repaired Message Handling

Messages are considered *repaired* if the message was completely scanned and all viruses have been repaired or removed. These messages will be delivered as is.

Encrypted Message Handling

Messages are considered *encrypted* if the engine is unable to finish the scan due to an encrypted or protected field in the message. Messages that are marked encrypted may also be repaired.

Note the differences between the encryption detection message filter rule (refer to "Encryption Detection Rule" in the "Using Message Filters to Enforce Email Policies" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*) and the virus scanning actions for "encrypted" messages. The encrypted message filter rule evaluates to "true" for any messages that are PGP or S/MIME encrypted. The encrypted rule can only detect PGP and S/MIME encrypted data. It does not detect password protected ZIP files, or Microsoft Word and Excel documents that include encrypted content. The virus scanning engine considers any message or attachment that is password protected to be "encrypted."

Note If you upgrade from a 3.8 or earlier version of AsyncOS and you configured Sophos Anti-Virus scanning, you must configure the Encrypted Message Handling section after you upgrade.

Unscannable Message Handling

Messages are considered *unscannable* if a scanning timeout value has been reached, or the engine becomes unavailable due to an internal error. Messages that are marked unscannable may also be repaired.

Virus Infected Message Handling

The system may be unable to drop the attachment or completely repair a message. In these cases, you can configure how the system handles messages that could still contain viruses.

The configuration options are the same for encrypted messages, unscannable messages, and virus messages.

Configuring Settings for Message Handling Actions

Action to Apply

Choose which overall action to take on each message type for encrypted, unscannable, or virus positive messages: drop the message, deliver the message as an attachment to a new message, deliver the message as is, or send the message to the anti-virus quarantine area (Quarantines and Anti-Virus Scanning, page 8-10). See the "Quarantines" chapter in *Cisco IronPort AsyncOS for Email Daily Management Guide* for more information about quarantines.

Configuring the appliance to deliver the infected messages as an attachment to a new message allows the recipient to choose how to deal with the original, infected attachment.

If you choose to deliver the message or deliver the message as an attachment to a new message, you can additionally:

- Modify message subject
- Archive original message
- Send generic notification
  The following actions are available in the "Advanced" section of the GUI:
- Add custom header to message
- Modify message recipient
- Send message to alternate destination host
- Send custom alert notification (to recipient only)

Note These actions are not mutually exclusive; you can combine some or all of them differently within different incoming or outgoing policies for different processing needs for groups of users. See the following sections and Notes on Anti-Virus Configurations, page 8-16 for more information on defining various scanning policies using these options.

Note Repaired messages have only two advanced options: Add custom header and Send custom alert notification. All other message types have access to all of the advanced options.

Quarantines and Anti-Virus Scanning

When flagged for quarantine, the message continues through the rest of the email pipeline. When the message reaches the end of the pipeline, if the message has been flagged for one or more quarantines then it enters those queues. Note that if the message does not reach the end of the pipeline, it is not placed in a quarantine.

For example, a content filter can cause a message to be dropped or bounced, in which case the message will not be quarantined.

Modify the Message Subject Header

You can alter the text of identified messages by prepending or appending certain text strings to help users more easily identify and sort identified messages.

Note White space is *not* ignored in the "Modify message subject" field. Add spaces after (if prepending) or before (if appending) the text you enter in this field to separate your added text from the original subject of the message. For example, add the text `[WARNING: VIRUS REMOVED]` with a few trailing spaces if you are prepending.

The default text is:

Table 8-1 Default Subject Line Text for Anti-Virus Subject Line Modification

Verdict	Default Text to Add to Subject
Encrypted	`[WARNING: MESSAGE ENCRYPTED]`
Infected	`[WARNING: VIRUS DETECTED]`
Repaired	`[WARNING: VIRUS REMOVED]`
Unscannable	`[WARNING: A/V UNSCANNABLE]`

Any message with multiple states causes a multi-part notification message informing users what actions the appliance performed on the message (for example, the user is notified that the message was repaired of a virus, but another part of the message was encrypted).

Archive Original Message

You can archive messages the system has identified as containing (or possibly containing) viruses to the "avarchive" directory. The format is an mbox-format log file. You *must* configure an "Anti-Virus Archive" log subscription to archive messages with viruses or messages that could not be completely scanned. For more information, refer to "Logging" in the *Cisco IronPort AsyncOS for Email Daily Management Guide* for more information.

Note In the GUI, you may need to click the "Advanced" link to reveal the "Archive original message" setting.

Sending Notifications

When the system has identified a message as containing viruses, you can send the default notification to the sender, the recipient, and/or additional users. When specifying additional users to notify, separate multiple addresses with commas (in both the CLI and the GUI). The default notification messages are:

Table 8-2 Default Notifications for Anti-Virus Notifications

Verdict	Notification
Repaired	The following virus(es) was detected in a mail message: <virus name(s)> Actions taken: Infected attachment dropped (or Infected attachment repaired).
Encrypted	The following message could not be fully scanned by the anti-virus engine due to encryption.
Unscannable	The following message could not be fully scanned by the anti-virus engine.
Infectious	The following unrepairable virus(es) was detected in a mail message: <virus name(s)>.

Add Custom Header to Message

You can define an additional, custom header to be added to all messages that are scanned by the anti-virus scanning engine. Click **Yes** and define the header name and text.

You can also create filters that use the `skip-viruscheck` action so that certain messages bypass virus scanning. See "Bypass Anti-Virus System Action" in the "Using Message Filters to Enforce Email Policies" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information.

Modify message recipient

You can modify the message recipient, causing the message to be delivered to a different address. Click **Yes** and enter the new recipient address.

Send message to alternate destination host

You can choose to send the notification to a different recipient or destination host for encrypted, unscannable, or virus infected messages. Click **Yes** and enter an alternate address or host.

For example, you could route suspected messages to an administrator's mailbox or a special mail server for subsequent examination. In the case of a multi-recipient message, only a single copy is sent to the alternative recipient.

Send custom alert notification (to recipient only)

You can send a custom notification to the recipient. To do so, you must first create the custom notification prior to configuring the settings. See Understanding Text Resources, page 14-12 for more information.

Figure 8-2 Options for Handling Messages Scanned for Viruses

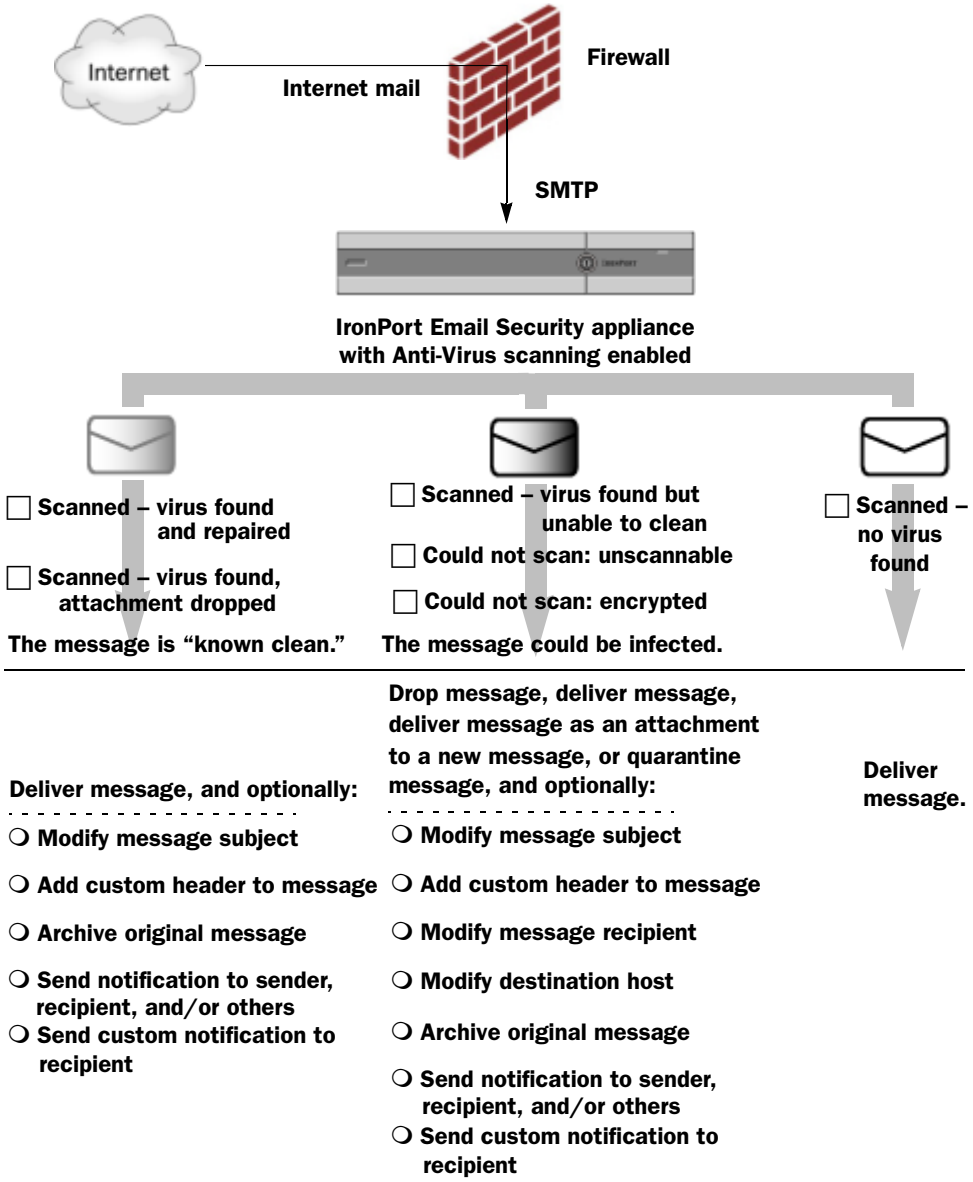

Note By default, Anti-Virus scanning is enabled in the $TRUSTED mail flow policy for public listeners, which is referenced by the WHITELIST sender group. See Mail Flow Policies: Access Rules and Parameters, page 5-8.

Editing the Anti-Virus Settings for a Mail Policy

The process for editing the per-user anti-virus settings for a mail policy is essentially the same for incoming or outgoing mail.

Individual policies (not the default) have an additional field to "Use Default" settings. Select this setting to inherit the default mail policy settings.

You enable anti-virus actions on a per-recipient basis using the Email Security Feature: the Mail Policies > Incoming or Outgoing Mail Policies pages (GUI) or the `policyconfig -> antivirus` command (CLI). After you enable anti-virus settings globally, you configure these actions separately for each mail policy you create. You can configure different actions for different mail policies.

Step 1 Click the link for the anti-virus security service in any row of the Email Security Manager incoming or outgoing mail policy table.

The Anti-Virus settings page similar to the one shown in Figure 8-3 and Figure 8-4 is displayed.

Click the link in the default row to edit the settings for the default policy. Figure 8-3 and Figure 8-4 show the settings for an individual policy (not the default).

Step 2 Click **Yes** or **Use Default** to enable Anti-Virus Scanning for the policy.

The first setting on the page defines whether the service is enabled for the policy. You can click **Disable** to disable the service altogether.

For mail policies other than the default, choosing "Yes" enables the fields in the Repaired, Encrypted, Unscannable, and Virus Infected Messages areas to become active.

Step 3 Select an Anti-Virus scanning engine. You can select McAfee or Sophos engines.

Step 4 Configure Message Scanning settings.

See Message Scanning Settings, page 8-8 for more information.

Step 5 Configure settings for Repaired, Encrypted, Unscannable, and Virus Infected messages.

Figure 8-3 and Figure 8-4 show the Anti-Virus settings for the mail policy named "Engineering" about to be edited. See Message Handling Settings, page 8-9 and Configuring Settings for Message Handling Actions, page 8-10.

Step 6 Click **Submit**.

The Mail Policies > Incoming or Outgoing Mail Policies page is refreshed to reflect the values you chose in the previous steps.

Step 7 Commit your changes.

Figure 8-3 Anti-Virus Settings for a Mail Policy (not default) - 1 of 2

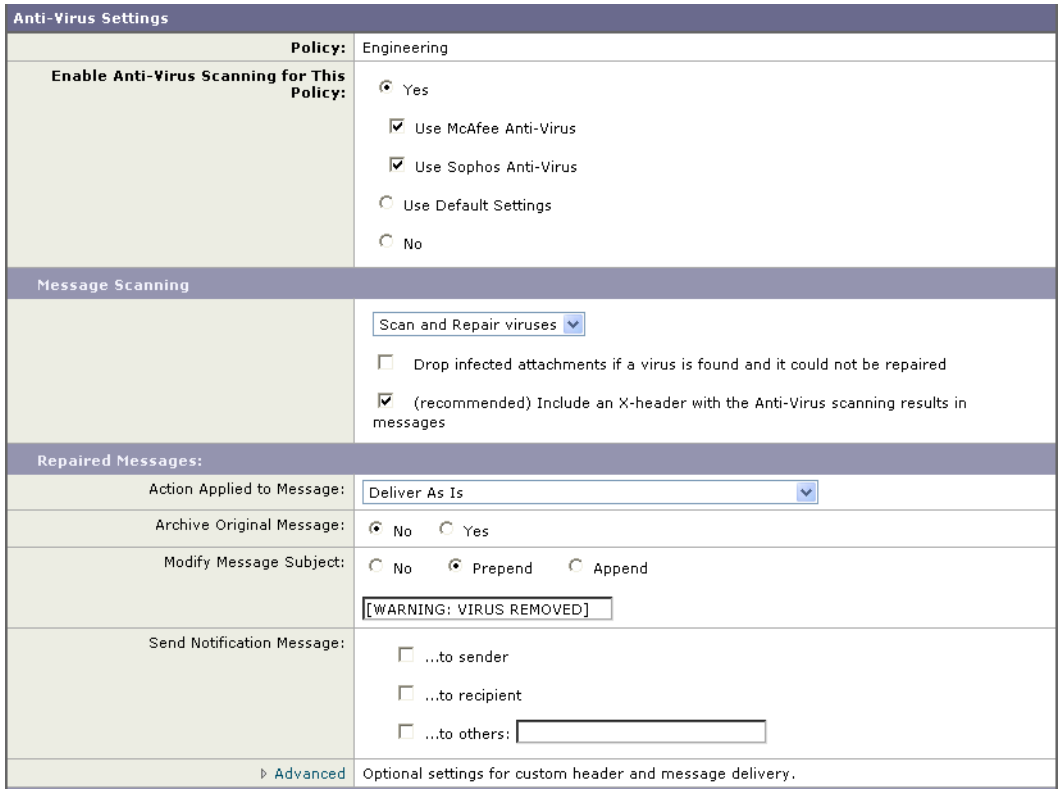

Figure 8-4 Anti-Virus Settings for a Mail Policy (not default) - 2 of 2

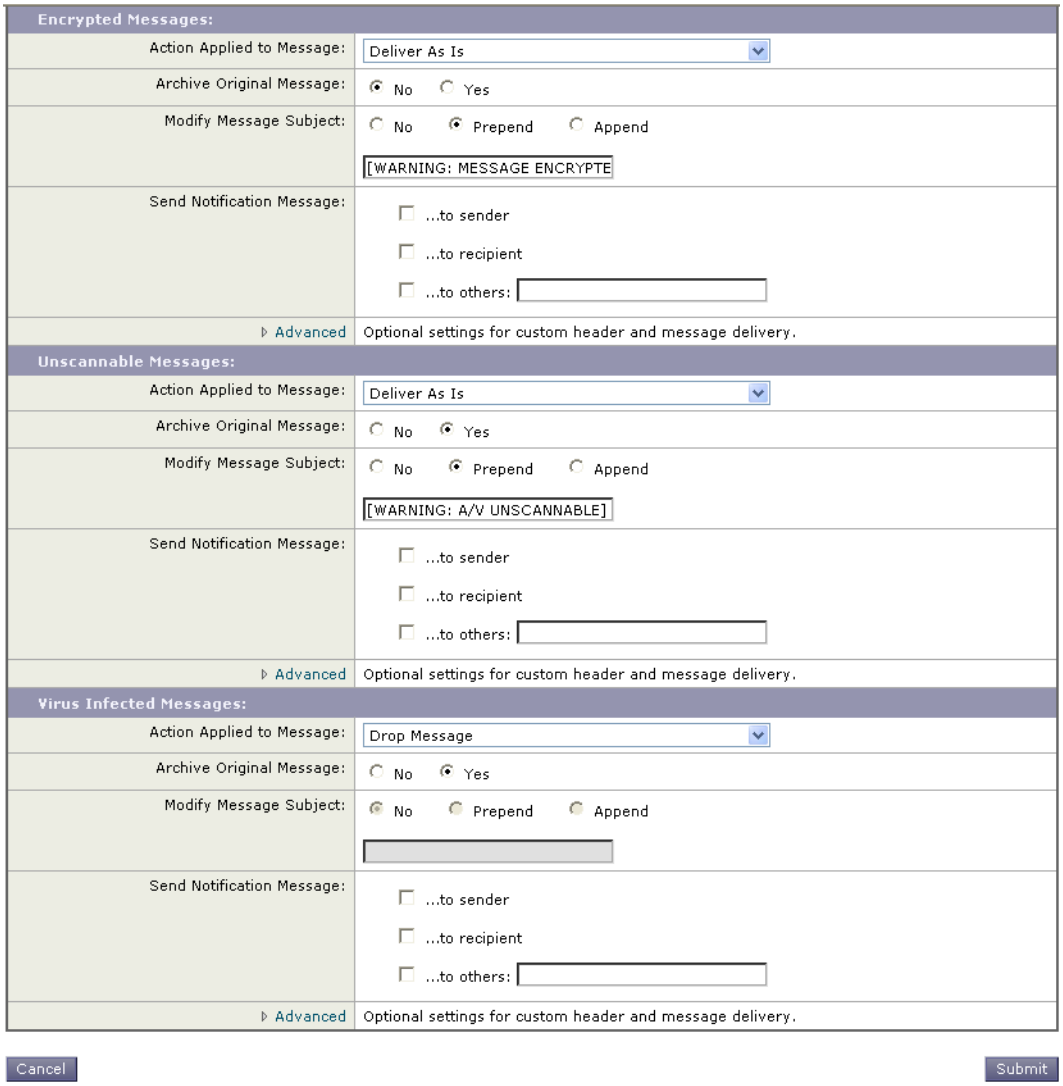

Notes on Anti-Virus Configurations

The drop attachments flag makes a considerable difference in how anti-virus scanning works. When the system is configured to "Drop infected attachments if a virus is found and it could not be repaired," any viral or unscannable MIME parts are removed from messages. The output from Anti-Virus scanning, then, is almost always a *clean* message. The action defined for *Unscannable Messages*, as shown in the GUI pane, rarely takes place.

In a "Scan for Viruses only" environment, these actions "clean" messages by dropping the bad message parts. Only if the RFC822 headers themselves are attacked or encounter some other problem would this result in the unscannable actions taking place. However, when Anti-Virus scanning is configured for "Scan for Viruses only" and "Drop infected attachments if a virus is found and it could not be repaired," is *not* chosen, the unscannable actions are very likely to take place.

Table 8-3 lists some common Anti-Virus configuration options.

Table 8-3 Common Anti-Virus Configuration Options

Situation	Anti-Virus Configuration
Widespread Virus Outbreak **Any viral message is simply dropped from the system with little other processing taking place.**	**Drop-attachments:** NO **Scanning:** Scan-Only **Cleaned messages**: Deliver **Unscannable messages**: DROP message **Encrypted messages:** Send to administrator or quarantine for review. **Viral messages:** Drop message
Liberal Policy **As many documents as possible are sent.**	**Drop-attachments:** YES **Scanning:** Scan and Repair **Cleaned messages:** [VIRUS REMOVED] and Deliver **Unscannable messages**: Forward as attachment **Encrypted messages:** Mark and forward **Viral messages**: Quarantine or mark and forward.
More Conservative Policy	**Drop-attachments**: YES **Scanning:** Scan and Repair **Cleaned messages:** [VIRUS REMOVED] and Deliver (Archive cleaned messages for a more cautious policy.) **Unscannable messages**: Send notification(s), quarantine, OR drop and archive. **Encrypted messages:** Mark and forward OR treat as unscannable **Viral messages:** Archive and drop
Conservative with Review **Possible virus messages are sent to a quarantine mailbox so that an administrator can review the content.**	**Drop-attachments:** NO **Scanning:** Scan-Only **Cleaned messages:** Deliver (this action won't normally be taken) **Unscannable messages**: Forward as attachment, `alt-src-host`, or `alt-rcpt-to` actions. **Encrypted messages:** Treat as unscannable **Viral messages:** Forward to quarantine or administrator.

Flow Diagram for Anti-Virus Actions

Figure 8-5 on page 8-18 explains how anti-virus actions and options affect messages processed by the appliance.

Figure 8-5 Flow Diagram for Anti-Virus Actions

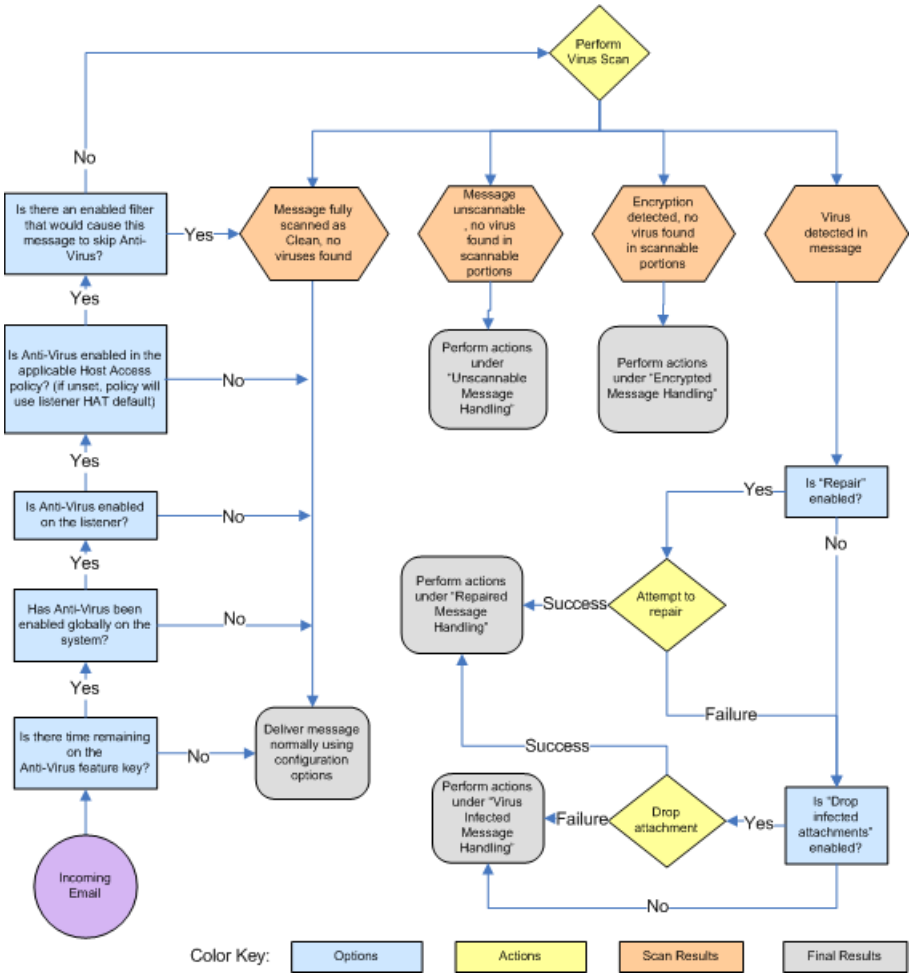

Note If you configure multi-layer anti-virus scanning, the Cisco IronPort appliance performs virus scanning with the McAfee engine first and the Sophos engine second. It scans messages using both engines, unless the McAfee engine detects a virus. If the McAfee engine detects a virus, the Cisco IronPort appliance performs the anti-virus actions (repairing, quarantining, etc.) defined for the mail policy.

Testing Virus Scanning

Step 1 Enable virus scanning for a mail policy.

Use the Security Services > Sophos/McAfee Anti-virus page or the `antivirusconfig` command to set global settings, and then use the Email Security Manager pages (GUI) or the `antivirus` subcommand of `policyconfig` to configure the settings for a specific mail policy.

Step 2 Open a standard text editor, then type the following character string as *one line, with no spaces or line breaks:*

```
X5O!P%@AP[4\PZX54(P^)7CC)7}$EICAR-STANDARD-ANTIVIRUS-TEST-FILE!$H+H*
```

Note The line shown above should appear as one line in your text editor window, so be sure to maximize your text editor window and delete any line breaks. Also, be sure to type the letter O, not the number 0, in the "X5O..." that begins the test message.

If you are reading this manual on your computer, you can copy the line directly from the PDF file or HTML file and paste it into your text editor. If you copy the line, be sure to delete any extra carriage returns or spaces.

Step 3 Save the file with the name **EICAR.COM**.

The file size will be 68 or 70 bytes.

Note This file is *not* a virus — it cannot spread or infect other files, or otherwise harm your computer. However, you should delete the file when you have finished testing your scanner to avoid alarming other users.

Step 4 Attach the file EICAR.COM to an email message, and send it to the listener that will match the mail policy you configured in step 1.

Ensure the that the recipient you specify in the test message will be accepted on the listener. (For more information, see Accepting Email for Local Domains or Specific Users on Public Listeners (RAT), page 5-50.)

Note that it may be difficult to email the file if you have virus scanning software is installed for outgoing mail on a gateway other than the Cisco IronPort (for example, a Microsoft Exchange server).

Note The test file always scans as unrepairable.

Step 5 Evaluate the actions you configured for virus scanning on the listener and ensure they are enabled and working as expected.

This is most easily accomplished by performing one of the following actions:

- Configure the virus scanning settings to Scan and Repair mode or Scan only mode without dropping attachments.

  Send an email with the Eicar test file as an attachment.

  Confirm that the actions taken match your configuration for Virus Infected Message Handling (the settings in Virus Infected Message Handling, page 8-10).

- Configure the virus scanning settings to Scan and Repair mode or Scan only mode with dropping attachments.

  Send an email with the Eicar test file as an attachment.

  Confirm that the actions taken match your configuration for Repaired Message Handling (the settings in Repaired Message Handling, page 8-9).

For more information obtaining virus files for testing anti-virus scanning, see:
`http://www.eicar.org/anti_virus_test_file.htm`
This page provides 4 files for downloading. Note that it may be difficult to download and extract these files if you have a client-side virus scanning software installed.

CHAPTER 9

Anti-Spam

The Cisco IronPort appliance offers a unique, layered approach to stopping spam at the email gateway. The first layer of spam control, reputation filtering (discussed previously in Chapter 7, "Reputation Filtering") allows you to classify email senders and restrict access to your email infrastructure based on senders' trustworthiness as determined by the Cisco IronPort SenderBase™ Reputation Service. The second layer of defense, scanning, is powered by Cisco IronPort Anti-Spam™ and Cisco IronPort Intelligent Multi-Scan technology. Coupled together, reputation filtering and anti-spam scanning offer the most effective and highest performing anti-spam solution available today.

Using the Cisco IronPort appliance, it is very easy to create policies to deliver messages from known or highly reputable senders — such as customers and partners — directly to the end user without any anti-spam scanning. Messages from unknown or less reputable senders can be subjected to anti-spam scanning, and you can also throttle the number of messages you are willing to accept from each sender. Email senders with the worst reputation can have their connections rejected or their messages dropped based on your preferences.

The unique, two-layer approach to fighting spam of the Cisco IronPort appliance provides you with a powerful and unprecedented flexibility to manage and protect your enterprise email gateway.

Anti-Spam Overview

Your Cisco IronPort appliance offers two anti-spam solutions: the Cisco IronPort Anti-Spam engine and Cisco IronPort Intelligent Multi-Scan. You can license and enable these solutions on your Cisco IronPort appliance, but you cannot enable both for the same policy. Using the Email Security Manager, you can quickly and easily specify a different anti-spam solution for different groups of users.

Enabling Anti-Spam Scanning

When using the System Setup Wizard (or `systemsetup` command in the CLI), you are presented with option to enable either Cisco IronPort Intelligent Multi-Scan or the Cisco IronPort Anti-Spam engine. You cannot enable both during system setup, but you can enable the anti-spam solution that you didn't choose by using the Security Services menu after system setup is complete. During system setup, you have the option to enable the Cisco IronPort Spam Quarantine for positive and suspect spam.

To enable the engine for the first time (either during system setup or later), read and agree to the license agreement.

Figure 9-1 Anti-Spam Engine - Selecting During System Setup

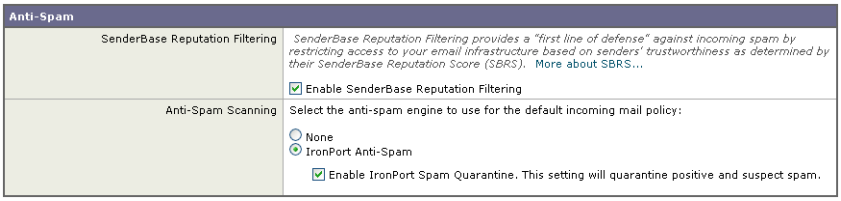

Note Please see Email Pipeline and Security Services, page 4-6 for information about how and when anti-spam scanning is applied.

After the system is set up, you can configure the anti-spam scanning solution for incoming mail policies via the Mail Policies > Incoming Mail Policies page. (Anti-spam scanning is typically disabled for outgoing mail policies.) You can even disable anti-spam scanning for a policy.

In this example, the default mail policy and the "Partners" policy are using the Cisco IronPort Anti-Spam scanning engine to quarantine positive and suspected spam.

Figure 9-2 Mail Policies - Anti-Spam Engine Per Recipient

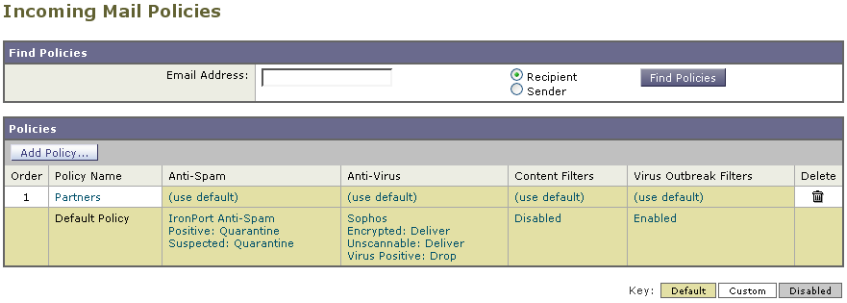

To change the Partners policy to use Cisco IronPort Intelligent Multi-Scan and scan for unwanted marketing messages, click on the entry in the Anti-Spam column corresponding with the Partners row ("use default").

Select Cisco IronPort Intelligent Multi-Scan for the scanning engine, and select Yes to enable unwanted marketing message detection. Use the default settings for unwanted marketing message detection.

Figure 9-3 shows Cisco IronPort Intelligent Multi-Scan and unwanted marketing message detection enabled in a policy.

Figure 9-3 Mail Policies - Enabling Cisco IronPort Intelligent Multi-Scan

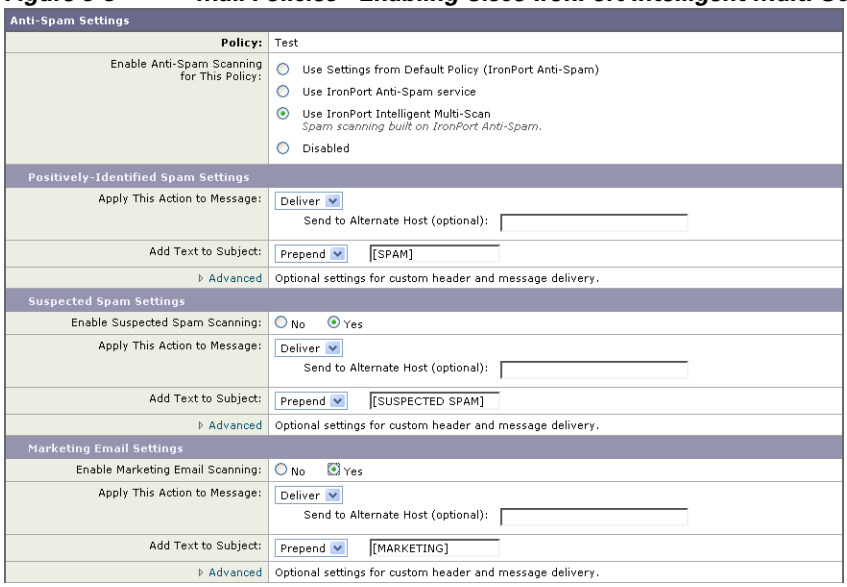

After submitting and committing the changes, the mail policy looks like this:

Figure 9-4 Mail Policies - Intelligent Multi-Scan Enabled in Policy

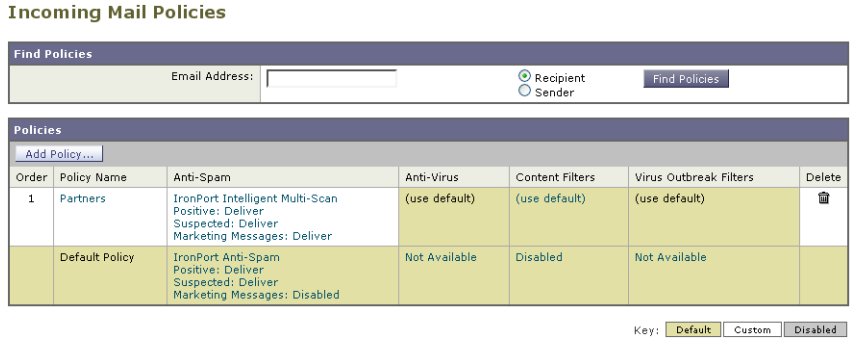

Anti-Spam Scanning Engine Settings

Each anti-spam solution has a group of configuration settings associated with it. These settings apply only to the corresponding engine, and are available on the Cisco IronPort Anti-Spam page and Cisco IronPort Intelligent Multi-Scan page on the Security Services menu, and the Incoming/Outgoing Mail Policies Anti-Spam settings page. The scanning solution's specific settings are discussed in the corresponding sections. The Cisco IronPort Anti-Spam and Cisco IronPort Intelligent Multi-Scan pages also display a list of their most recent anti-spam rule updates.

For more information on configuring global anti-spam settings, see:

- Enabling Cisco IronPort Anti-Spam and Configuring Global Settings, page 9-6 and
- Enabling Cisco IronPort Intelligent Multi-Scan and Configuring Global Settings, page 9-9.

For more information on configuring anti-spam scanning on a per recipient basis, see Configuring Per-Recipient Policies for Anti-Spam, page 9-12.

Anti-Spam Scanning and Messages Generated by the Cisco IronPort Appliance

Cisco recommends that recipients who receive email alerts, scheduled reports, and other automated messages from the Cisco IronPort appliance be placed in an incoming mail policy that bypasses anti-spam scanning. These messages may contain URLs or other information associated with spam sources not ordinarily found in a company's mail stream which may occasionally cause such messages to be marked as SPAM. Alternatively, you can choose to add the IP addresses sending mail on behalf of the Cisco IronPort appliance to the 'WHITELIST' policy in the host access table (see Adding a Sender to a Sender Group, page 5-33). For more information, please contact your authorized Cisco IronPort appliance support center.

Cisco IronPort Anti-Spam Filtering

Cisco IronPort Anti-Spam uses conventional techniques and innovative context-sensitive detection technology to eliminate a diverse range of known and emerging email threats.

Evaluation Key

Your Cisco IronPort appliance ships with a 30-day evaluation key for the Cisco IronPort Anti-Spam software. This key is not enabled until you accept the license agreement in the system setup wizard or Security Services > IronPort Anti-Spam pages (in the GUI) or the `systemsetup` or `antispamconfig` commands (in the CLI). Once you have accepted the agreement, Cisco IronPort Anti-Spam will be enabled, by default, for the default incoming Mail Policy. An alert is also sent to the administrator address you configured (see Step 2: System, page 3-15) noting that the Cisco IronPort Anti-Spam license will expire in 30 days. Alerts are sent 30, 15, 5, and 0 days prior to expiration. For information on enabling the feature beyond the 30-day evaluation period, contact your Cisco IronPort sales representative. You can see how much time remains on the evaluation via the System Administration > Feature Keys page or by issuing the `featurekey` command. (For more information, see the section on working with feature keys in “Common Administrative Tasks” in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.)

Cisco IronPort Anti-Spam and CASE: an Overview

Cisco IronPort Anti-Spam filtering is based on Context Adaptive Scanning Engine (CASE) ™, and is the first anti-spam scanning engine to combine email and web reputation information to:

- Eliminate the broadest range of email threats — detect spam, “phishing,” zombie-based attacks, and other “blended” threats.
- Deliver the highest accuracy — anti-spam rules based on email and web reputation from SenderBase Reputation Service.
- Offer ease of use — due to reduced hardware and administrative costs.
- Deliver industry leading performance — CASE uses dynamic early exit criteria and off-box network calculations to deliver breakthrough performance.
- Address the needs of international users — Cisco IronPort Anti-Spam is tuned to deliver industry-leading efficacy world-wide.

Broadest Threat Prevention

CASE combines content analysis, email reputation, and web reputation to deliver the broadest set of threat prevention factors.

Cisco designed Cisco IronPort Anti-Spam from the ground up to detect the broadest range of email threats. Cisco IronPort Anti-Spam addresses a full range of known threats including spam, phishing and zombie attacks, as well as hard-to-detect low volume, short-lived email threats such as “419” scams. In addition, Cisco IronPort Anti-Spam identifies new and evolving blended threats such as spam attacks distributing malicious content through a download URL or an executable.

To identify these threats, Cisco IronPort Anti-Spam uses the industry's most complete approach to threat detection, examining the full context of a message-its content, methods of message construction, the reputation of the sender, and the reputation of web sites advertised in the message and more. Only Cisco IronPort Anti-Spam combines the power of email and web reputation data, leveraging the full power of the world's largest email and web traffic monitoring network — SenderBase — to detect new attacks as soon as they begin.

Note If your Cisco IronPort appliance is set to receive mail from a local MX/MTA, you must identify upstream hosts that may mask the sender’s IP address. See Incoming Relays, page 9-19 for more information.

Lowest False Positive Rate

Cisco IronPort Anti-Spam and Cisco IronPort Outbreak Filters are powered by Cisco IronPort’s patent-pending Context Adaptive Scanning Engine (CASE) ™. CASE provides breakthrough accuracy and performance by analyzing over 100,000 message attributes across four dimensions:

Step 1 Email reputation — *who* is sending you this message?

Step 2 Message content — *what* content is included in this message?

Step 3 Message structure — *how* was this message constructed?

Step 4 Web reputation — *where* does the call to action take you?

Analyzing multi-dimensional relationships allows CASE to catch a broad range of threats while maintaining exceptional accuracy. For example, a message that has content claiming to be from a legitimate financial institution but that is sent from an IP address on a consumer broadband network or that contains a URL hosted on a “zombie” PC will be viewed as suspicious. In contrast, a message coming from a pharmaceutical company with a positive reputation will not be tagged as spam even if the message contains words closely correlated with spam.

Industry-Leading Performance

CASE combines the following features to deliver accurate verdicts quickly:

- Multiple threats are scanned for in a single pass
- Dynamic “early exit” system

  System performance is optimized using Cisco IronPort's unique “early exit” system. Cisco IronPort developed a proprietary algorithm to determine the order in which rules are applied based on rule accuracy and computational expense. Lighter and more accurate rules are run first, and if a verdict is reached, additional rules are not required. This improves system throughput, allowing our

products to meet the needs of large-scale enterprises. Conversely, the efficiency of the engine allows for implementation on low-cost hardware, making Cisco IronPort's security services attractive for low-end customers.

- Off-box network calculations

International Users

Cisco IronPort Anti-Spam is tuned to deliver industry-leading efficacy world-wide. In addition to locale-specific content-aware threat detection techniques, you can further optimize anti-spam scanning for specific regions using regional rules profiles. The anti-spam engine includes a regional rules profile. The regional rules profile targets spam on a regional basis. For example, China and Taiwan receive a high percentage of spam in traditional or modern Chinese. The Chinese regional rules are optimized for this type of spam. Cisco strongly recommends you use the Chinese regional rules profile if you receive mail primarily for mainland China, Taiwan, and Hong Kong. You can enable the regional rules profile from Security Services > IronPort Anti-Spam.

Note Because the regional rules profile optimizes the anti-spam engine for a particular region, it can reduce capture rates for other types of spam. Therefore, Cisco recommends you enable this feature only if you receive the bulk of your email from the specified region.

Cisco IronPort Anti-Spam leverages globally representative email and web content-agnostic data contributed by over 125,000 ISPs, universities and corporations throughout the Americas, Europe, and Asia. The Threat Operations Center is set up for global operations with centers in Sao Paulo, Beijing and London. In addition, analysts speak 32 languages including Chinese, Japanese, Korean, Portuguese, and Spanish.

Enabling Cisco IronPort Anti-Spam and Configuring Global Settings

Overview

You enable Cisco IronPort Anti-Spam and modify its global configuration settings using the Security Services > IronPort Anti-Spam and Security Services > Service Updates pages (GUI) or the `antispamconfig` and `updateconfig` commands (CLI). The following global settings are configured:

- Enable Cisco IronPort Anti-Spam globally for the appliance.
- Configure the thresholds for message scanning by Cisco IronPort Anti-Spam.

  To optimize the throughput of your appliance while still being able to scan the increasing larger messages sent by spammers, you can define an *always scan* message size, where messages smaller than the defined size are completely scanned by CASE, delivering Cisco IronPort's industry-leading level of efficacy, and a *never scan* message size, where messages larger than the defined size are not scanned by CASE. For messages larger than the *always scan* size and smaller than the *never scan* size, CASE performs a limited and faster scan.

  Note If the Outbreak Filters maximum message size is greater than Cisco IronPort Anti-Spam's *always scan* message, CASE fully scans messages smaller than the Outbreak Filters maximum size.

- Enter a length of time to wait for timeout when scanning a message.

- Define and (optionally) enable a proxy server for obtaining Cisco IronPort Anti-Spam rules updates (Security Services > Service Updates). If you define a proxy server to retrieve rules updates, you can optionally configure an authenticated username, password, and specific port when connecting to the proxy server.
- Define and (optionally) enable a download server from which to receive Cisco IronPort Anti-Spam rules updates (Security Services > Service Updates).
- Enable or disable receiving automatic updates to Cisco IronPort Anti-Spam rules, and also specify the update interval.

Note The proxy server setup is available via the Security Services > Service Updates page. For more information about specifying a proxy server, see The Service Updates Page, page 15-10. Note that the proxy server is global in that all services that are configured to use a proxy server will use the same proxy server.

Note If you chose to enable Cisco IronPort Anti-Spam in the GUI's system setup wizard (or the CLI `systemsetup` command), it will be enabled for the default incoming mail policy with the default values for the global settings.

Figure 9-5 shows the global settings that you configure on the Security Services > IronPort Anti-Spam page.

Figure 9-5 Cisco IronPort Anti-Spam Global Settings: Editing

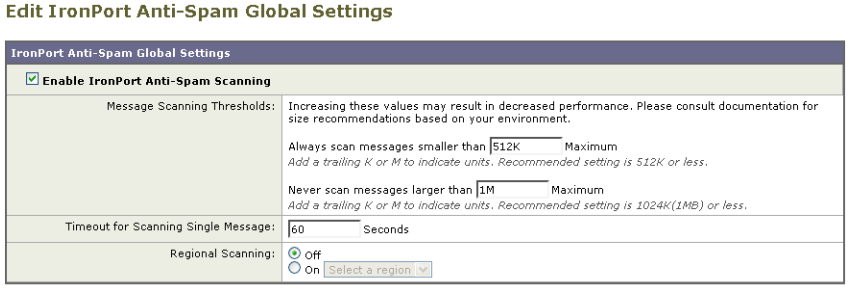

Step 1 If you have not enabled Cisco IronPort Anti-Spam in the system setup wizard, select Security Services > IronPort Anti-Spam.

Step 2 Click **Enable**.

The license agreement page is displayed.

Note If you do not accept the license agreement, Cisco IronPort Anti-Spam is not enabled on the appliance.

Step 3 Scroll to the bottom of the page and click **Accept** to accept the agreement.

A page similar to Figure 9-6 is displayed.

Step 4 Click **Edit Global Settings**.

Step 5 Check the box next to Enable IronPort Anti-Spam scanning.

Checking this box enables the feature globally for the appliance. However, you must still enable per-recipient settings in Mail Policies. For more information, see Configuring Per-Recipient Policies for Anti-Spam, page 9-12

Step 6 Enter a value for the *always scan* message size for Cisco IronPort Anti-Spam.

The recommended value is 512 Kb or less. Messages smaller than the *always scan* size will be fully scanned by CASE, except in cases of "early exit." Messages larger than this size are partially scanned by CASE if they are smaller than the *never scan* size entered in Step 7. See Industry-Leading Performance, page 9-5 for more information on the "early exit" system.

Note Cisco advises not to exceed 3 MB for the *always scan* message size. A larger value may result in decreased performance.

Step 7 Enter a value for the *never scan* message size.

The recommended value is 1024 Kb or less. Messages larger than this size will not be scanned by Cisco IronPort Anti-Spam and the `X-IronPort-Anti-Spam-Filtered: true` header will not be added to the message.

Note Cisco advises not to exceed 10 MB for the *never scan* message size. A larger value may result in decreased performance.

Step 8 Enter the number of seconds to wait for timeout when scanning a message.

When specifying the number of seconds, enter an integer from 1 to 120. The default value is 60 seconds.

Step 9 Enable or disable regional scanning. Regional scanning optimizes Cisco IronPort Anti-Spam scanning for a particular region. Because this feature optimizes the anti-spam engine for a particular region, it can reduce capture rates for other types of spam. Therefore, Cisco recommends you enable this feature only if you receive the bulk of your email from the specified region. For more information about regional scanning, see International Users, page 9-6.

Step 10 Submit and commit your changes.

The Security Services > IronPort Anti-Spam page is refreshed to display the values you chose in the previous steps.

Figure 9-6 Cisco IronPort Anti-Spam Global Settings

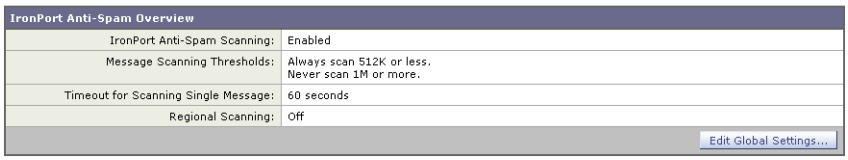

Additional Steps

Once you have enabled Cisco IronPort Anti-Spam, enable SenderBase Reputation Service scoring, even if you are not rejecting connections based on SenderBase Reputation Scores. For more information on enabling SBRS, see Implementing SenderBase Reputation Filters, page 7-4.

Cisco IronPort Intelligent Multi-Scan Filtering

Cisco IronPort Intelligent Multi-Scan incorporates multiple anti-spam scanning engines, including Cisco IronPort Anti-Spam, to provide an intelligent, multi-layer anti-spam solution. This method provides more accurate verdicts that increase the amount of spam that is caught but without increasing the false positives rate.

When processed by Cisco IronPort Intelligent Multi-Scan, a message is first scanned by third-party anti-spam engines. Cisco IronPort Intelligent Multi-Scan then passes the message and the verdicts of the third-party engines to Cisco IronPort Anti-Spam, which assumes responsibility for the final verdict. After Cisco IronPort Anti-Spam performs its scan, it returns a combined multi-scan score to AsyncOS. Combining the benefits of the third-party scanning engines and Cisco IronPort Anti-Spam results in more caught spam while maintaining Cisco IronPort Anti-Spam's low false positive rate.

You cannot configure the order of the scanning engines used in Cisco IronPort Intelligent Multi-Scan; Cisco IronPort Anti-Spam will always be the last to scan a message and Cisco IronPort Intelligent Multi-Scan will not skip it if a third-party engine determines that a message is spam.

Using Cisco IronPort Intelligent Multi-Scan can lead to reduced system throughput. Please contact your Cisco IronPort support representative for more information.

This feature is supported on all C-Series and X-Series appliances, except for the C100 appliance.

Note The Intelligent Multi-Scan feature key also enables Cisco IronPort Anti-Spam on the appliance, giving you the option of enabling either Cisco IronPort Intelligent MultiScan or Cisco IronPort Anti-Spam for a mail policy.

Enabling Cisco IronPort Intelligent Multi-Scan and Configuring Global Settings

Overview

You enable Cisco IronPortIntelligent Multi-Scan and modify its global configuration settings using the Security Services > IronPort Intelligent Multi-Scan and Security Services > Service Updates pages (GUI) or the `antispamconfig` and `updateconfig` commands (CLI). The following global settings are configured:

- Enable Cisco IronPort Intelligent Multi-Scan globally for the appliance.
- Configure the maximum size of message to be scanned by Cisco IronPort Intelligent Multi-Scan.
- Enter a length of time to wait for timeout when scanning a message.

  Most users will not need to change the maximum message size to be scanned or the timeout value. That said, you may be able to optimize the throughput of your appliance by lowering the maximum message size setting.
- Define and (optionally) enable a proxy server for obtaining Cisco IronPort Intelligent Multi-Scan rules updates (Security Services > Service Updates). If you define a proxy server to retrieve rules updates, you can optionally configure an authenticated username, password, and specific port when connecting to the proxy server.
- Define and (optionally) enable a download server from which to receive Cisco IronPort Intelligent Multi-Scan rules updates (Security Services > Service Updates).
- Enable or disable receiving automatic updates to Cisco IronPort Intelligent Multi-Scan rules, and also specify the update interval.

Note The proxy server setup is available via the Security Services > Service Updates page. For more information about specifying a proxy server, see The Service Updates Page, page 15-10. Note that the proxy server is global in that all services that are configured to use a proxy server will use the same proxy server.

Note If you chose to enable Cisco IronPort Intelligent Multi-Scan in the GUI's system setup wizard (or the CLI `systemsetup` command), it will be enabled for the default incoming mail policy with the default values for the global settings.

Figure 9-7 shows the global settings that you configure on the Security-Services > IronPort Intelligent Multi-Scan page.

Figure 9-7 Cisco IronPort Intelligent Multi-Scan Global Settings: Editing

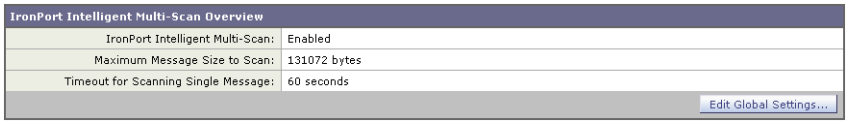

To enable Cisco IronPort Intelligent Multi-Scan, follow these steps:

Step 1 If you did not enable Cisco IronPort Intelligent Multi-Scan in the system setup wizard, select Security Services > IronPort Intelligent Multi-Scan.

Step 2 Click **Enable**.

The license agreement page is displayed.

Note If you do not accept the license agreement, Cisco IronPort Intelligent Multi-Scan is not enabled on the appliance.

Step 3 Scroll to the bottom of the page and click **Accept** to accept the agreement.

A page similar to Figure 9-8 is displayed.

Step 4 Click **Edit Global Settings**.

Step 5 Check the box next to Enable IronPort Intelligent Multi-Scan.

Checking this box enables the feature globally for the appliance. However, you must still enable per-recipient settings in Mail Policies. For more information, see Configuring Per-Recipient Policies for Anti-Spam, page 9-12.

Step 6 Choose a value for the *maximum message size to scan* by Cisco IronPort Intelligent Multi-Scan.

The default value is 128 Kb. Messages larger than this size will not be scanned by Cisco IronPort Intelligent Multi-Scan.

Step 7 Enter the number of seconds to wait for timeout when scanning a message.

When specifying the number of seconds, enter an integer from 1 to 120. The default value is 60 seconds.

Step 8 Submit and commit your changes.

The Security Services > IronPort Intelligent Multi-Scan page is refreshed to display the values you chose in the previous steps.

Figure 9-8 Cisco IronPort Intelligent Multi-Scan Global Settings

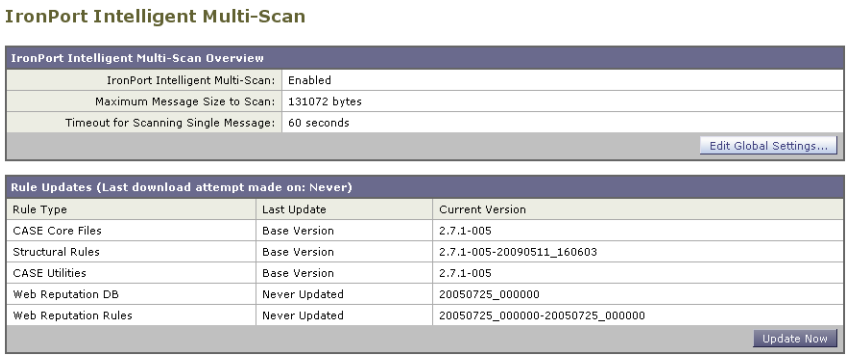

Additional Steps

Once you have enabled Cisco IronPort Intelligent Multi-Scan, enable SenderBase Reputation Service scoring, even if you are not rejecting connections based on SenderBase Reputation scores. For more information on enabling SBRS, see Implementing SenderBase Reputation Filters, page 7-4.

Configuring Anti-Spam Rule Updating

Cisco IronPort Anti-Spam and Cisco IronPort Intelligent Multi-Scan rules are retrieved (by default) from Cisco IronPort's update servers. You can specify a local server for updates, a proxy server to use for retrieving updates, and whether and how frequently to check for rule updates. To configure updates for your anti-spam solution, click **Edit Update Settings** on the Security Services > Service Updates page.

See Service Updates, page 15-10 for more information.

Enabling a Proxy Server for Obtaining Cisco IronPort Anti-Spam Rules Updates

The Cisco IronPort appliance is configured to connect directly to Cisco IronPort's update servers to receive anti-spam rules updates. This connection is made by HTTP on port 80 and the content is encrypted. If you do not want to open this port in your firewall, you can define a proxy server and specific port from which the appliance can receive updated rules.

If you choose to use a proxy server, you can specify an optional authentication and port.

Cisco IronPort Anti-Spam and Cisco IronPort Intelligent Multi-Scan will *automatically* use a proxy server if one has been defined. There is no way to turn off the proxy server for the anti-spam solution without disabling it for all other service updates (Outbreak Filters, Sophos Anti-Virus, etc.).

Note If you define a proxy server, it will be used for all service updates that are configured to use a proxy server, automatically.

For more information about defining a proxy server, see Specify an HTTP Proxy Server (Optional), page 15-14.

Monitoring Rules Updates

Once you have accepted the license agreement, the most recent Cisco IronPort Anti-Spam and Cisco IronPort Intelligent Multi-Scan rules updates are listed on the their corresponding page in the Security Services menu (GUI) and in the `antispamstatus` command (CLI).

Note If the update has not occurred, or a server has not been configured, the string "Never Updated" is displayed.

Figure 9-9 Rules Updates Section of Security Services > IronPort Anti-Spam Page: GUI

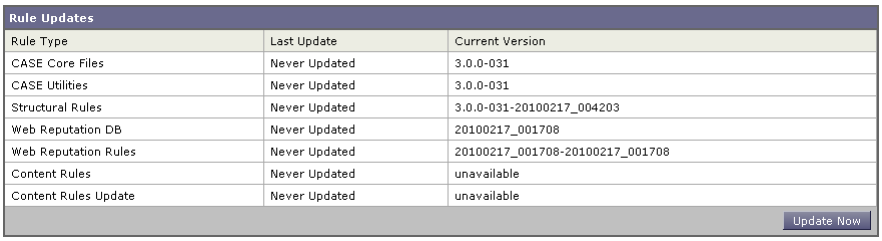

Configuring Per-Recipient Policies for Anti-Spam

The Cisco IronPort Anti-Spam and Cisco IronPort Intelligent Multi-Scan solutions process email for incoming (and outgoing) mail based on policies (configuration options) you configure using the Email Security Manager feature. Cisco IronPort Anti-Spam and Cisco IronPort Intelligent Multi-Scan scan messages through their filtering modules for classification. The classification, or *verdict*, is then returned for subsequent delivery action. Four verdicts are possible: messages can be identified as not spam, identified as a unwanted marketing email, positively identified as spam, or suspected to be spam. Actions taken on messages positively identified as spam, suspected to be spam, or identified as unwanted marketing messages include:

- Specifying a Positive or Suspected Spam threshold.
- Choosing which overall action to take on unwanted marketing messages, positively identified spam, or suspected spam messages: deliver, drop, bounce, or quarantine.
- Archiving messages to an mbox-format log file. You must create a log to enable archiving messages identified as spam. See Archiving Identified Messages, page 9-14.
- Altering the subject header of messages identified as spam or marketing.
- Sending messages to an alternate destination mailhost.
- Adding a custom X-Header to messages.
- Sending messages to an alternate envelope recipient address. (For example, you could route messages identified as spam to an administrator's mailbox for subsequent examination.) In the case of a multi-recipient message, only a single copy is sent to the alternate recipient.

Note These actions are not mutually exclusive; you can combine some or all of them differently within different incoming or outgoing policies for different processing needs for groups of users. You can also treat positively identified spam differently from suspected spam in the same policy. For example, you may want to drop messages positively identified as spam, but quarantine suspected spam messages.

You enable Cisco IronPort Anti-Spam or Cisco IronPort Intelligent Multi-Scan actions on a per-recipient basis using the Email Security Manager feature: the Mail Policies > Incoming or Outgoing Mail Policies pages (GUI) or the `policyconfig -> antispam` command (CLI). After the anti-spam solution has been enabled globally, you configure these actions separately for each mail policy you create. You can configure different actions for different mail policies.You can only enable one anti-spam solution per policy; you cannot enable both on the same policy.

Note To enable anti-spam scanning for outgoing mail, you also need to check the anti-spam settings of the relevant host access table, especially for a private listener. For more information, see Mail Flow Policies: Access Rules and Parameters, page 5-8.

Each row in the Email Security Manager represents a different policy. Each column represents a different security service.

Figure 9-10 Mail Policies - Anti-Spam Engine

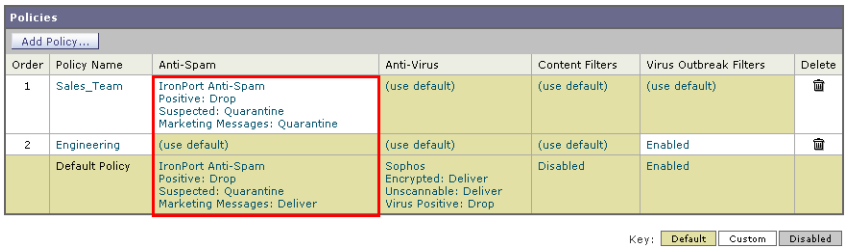

Editing the Anti-Spam Settings for a Mail Policy

The process for editing the per-user anti-spam settings for a mail policy is essentially the same, whether the policy is for incoming or outgoing mail.

Individual policies (not the default) have an additional field to "Use Default" settings. Selecting this causes the policy to adopt all of the Anti-Spam settings from the default mail policy.

See also Editing the Default Policy: Anti-Spam Settings, page 6-21 for more information.

Step 1 Click the link for the Anti-Spam security service in any row of the Email Security Manager incoming or outgoing mail policy table.

The Anti-Spam settings page similar to the one shown in Figure 9-11 is displayed.

Click the link in the default row to edit the settings for the default policy. Figure 9-11 shows the settings for a specific policy (not the default). Compare this screen with Figure 6-6 on page 6-22. Note how individual policies have the "Use Default" option.

Step 2 Select the anti-spam solution you want to use for the policy.

You can click **Disabled** to disable anti-spam scanning altogether for the mail policy.

Step 3 Configure settings for positively identified spam, suspected spam, and unwanted marketing messages.

Figure 9-11 shows the Cisco IronPort Anti-Spam settings for the default mail policy about to be edited. See Positively Identified versus Suspected Spam, page 9-16 and Notes on Configuring Settings for Identified Messages, page 9-14.

Step 4 Submit and commit your changes.

The Mail Policies > Incoming or Outgoing Mail Policies page is refreshed to reflect the values you chose in the previous steps.

Notes on Configuring Settings for Identified Messages

Positive/Suspected Spam Threshold

Enter a threshold value for positively identified spam and a value for suspected spam. For more information about spam thresholds, see Positive and Suspect Spam Threshold, page 9-15.

Action to Apply

Choose which overall action to take on positively identified spam, suspected spam, or unwanted marketing messages: Deliver, Drop, Bounce, or Quarantine.

Archiving Identified Messages

You can archive identified messages into the "Anti-Spam Archive" log. The format is an mbox-format log file. For more information, see the example below and refer to the "Logging" chapter of the *Cisco IronPort AsyncOS for Email Daily Management Guide*.

Altering the Subject Header

You can alter the text of the Subject header on identified messages by prepending or appending certain text strings to help users more easily identify and sort spam and unwanted marketing messages.

Note White space is *not* ignored in the "Modify message subject" field. Add spaces after (if prepending) or before (if appending) the text you enter in this field to separate your added text from the original subject of the message. For example, add the text `[SPAM]` with a few trailing spaces if you are prepending.

Note The "Add text to message" field only accepts US-ASCII characters.

Sending Identified Messages to an Alternate Destination Host

You can send identified messages to an alternate destination mailhost.

Adding a Custom X-Header

You can add a custom X-Header to identified messages.

Click **Yes** and define the header name and text.

Changing the Envelope Recipient Address

You can have identified messages sent to an alternate envelope recipient address.

Click **Yes** and define an alternate address.

For example, you could route messages identified as spam to an administrator's mailbox for subsequent examination. In the case of a multi-recipient message, only a single copy is sent to the alternate recipient.

Figure 9-11 Cisco IronPort Anti-Spam Settings for a Mail Policy

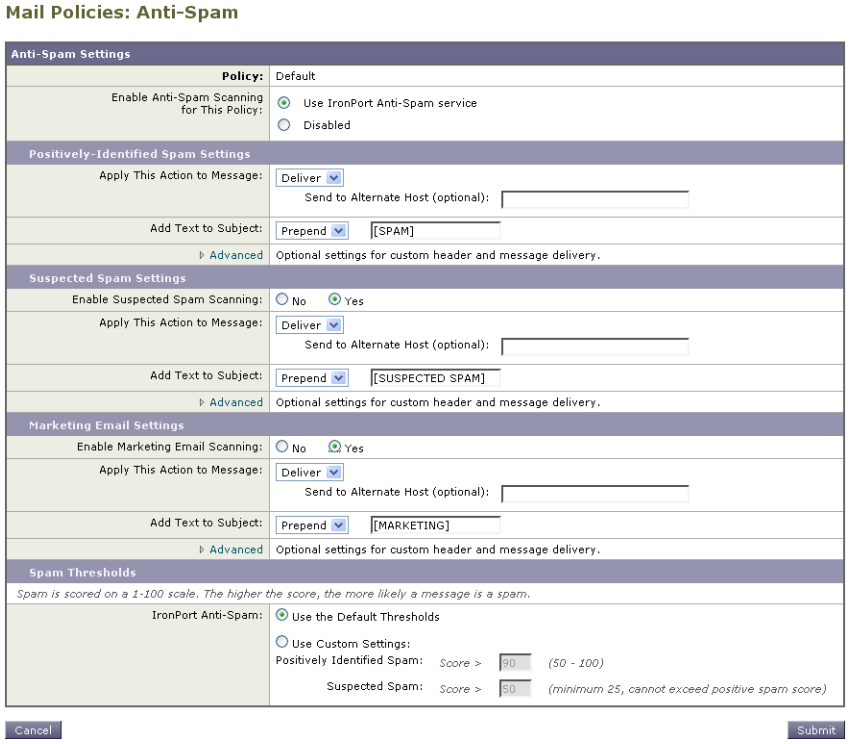

Positive and Suspect Spam Threshold

When evaluating messages for spam, Cisco IronPort Anti-Spam and Cisco IronPort Intelligent Multi-Scan apply thousands of rules in order to arrive at an overall spam score for the message. To maintain its high accuracy, the both anti-spam solutions by default set this threshold value quite high. Messages returning a score between 90 and 100 are considered to be positively identified as spam. You can change the positively identified spam threshold to a value between 75 (most aggressive) and 99 (most conservative). You can configure the anti-spam solution to reflect the spam tolerance levels of your organization. Both Cisco IronPort Anti-Spam and Cisco IronPort Intelligent Multi-Scan provide a configurable Positive and Suspected spam threshold, applicable *per mail policy*. This allows you to create an optional category of "suspected spam" — a gray area of messages that are suspiciously similar to spam, but also share some traits with legitimate messages.

You can change the threshold setting of this new category to different levels of aggressiveness, so that any messages with scores below the configured suspected spam range will be considered legitimate, and any messages above the suspected threshold but below the positive threshold will be considered to be suspected spam and will be treated accordingly. You can also define a separate action to take on suspected spam; for example, you may wish to drop "positively identified" spam, but quarantine "suspected" spam.

The higher the number you enter, the higher the threshold for Cisco IronPort Anti-Spam rules used to determine if a message qualifies as suspected spam. Enter a lower number to enable a lower threshold and subsequently mark more messages as "possible spam" (which may result in a higher false positive rate). Conversely, enter a higher number if you want to ensure that only spam messages are being filtered (which may result in some spam getting through). The default value is 50. See Positively Identified versus Suspected Spam, page 9-16 for common configurations using this two categories.

The suspected spam threshold is set per mail policy for Cisco IronPort Anti-Spam.

Positively Identified versus Suspected Spam

Because Cisco IronPort Anti-Spam and Cisco IronPort Intelligent Multi-Scan make the distinction between positively identified and suspected spam (Positive and Suspect Spam Threshold, page 9-15), many users configure their systems in one of the following ways:

Table 9-1 Common Example Configurations of Positively Identified and Suspected Spam

Spam	Method 1 Actions (Aggressive)	Method 2 Actions (Conservative)
Positively Identified	Drop	Deliver with "`[Positive Spam]`" added to the subject of messages
Suspected	Deliver with "`[Suspected Spam]`" added to the subject of messages	Deliver with "`[Suspected Spam]`" added to the subject of messages

The first configuration method tags only suspected spam messages, while dropping those messages that are positively identified. Administrators and end-users can check the subject line of incoming message for false positives, and an administrator can adjust, if necessary, the suspected spam threshold.

In the second configuration method, positively identified and suspected spam is delivered with an altered subject. Users can delete suspected and positively identified spam. This method is more conservative than the first.

See Table 6-6 on page 6-29 for a further discussion of mixing aggressive and conservative policies on a per-recipient basis using the Email Security Manager feature.

Unwanted Marketing Message Detection

Cisco IronPort Anti-Spam and Cisco IronPort Intelligent Multi-Scan can distinguish between spam and unwanted marketing messages from a legitimate source. Even though marketing messages are not considered spam, your organization or end-users may not want to receive them. Like spam, you have the option to deliver, drop, quarantine, or bounce unwanted marketing message. You also have the option to tag unwanted marketing messages by adding text to the message's subject to identify it as marketing.

Headers Added by Cisco IronPort Anti-Spam and Intelligent Multi-Scan

If Cisco IronPort Anti-Spam scanning or Intelligent Multi-Scan is enabled for a mail policy, each message that passes through that policy will have the following header added to the message:

```
X-IronPort-Anti-Spam-Filtered: true
```

A second header will also be inserted for each message filtered by Cisco IronPort Anti-Spam or Intelligent Multi-Scan. This header contains information that allows Cisco IronPort Support to identify the CASE rules and engine version used to scan the message:

```
X-IronPort-Anti-Spam: result
```

Cisco IronPort Intelligent Multi-Scan also adds headers from the third-party anti-spam scanning engines.

In addition, using the Email Security Manager feature, you can define an additional custom header to be added to all messages for a given policy that are positively identified as spam, suspected to be spam, or identified as unwanted marketing mail. (See Adding a Custom X-Header, page 9-14.)

You can also create message filters that use the `skip-spamcheck` action so that certain messages skip Cisco IronPort Anti-Spam scanning. For more information, refer to "Bypass Anti-Spam System Action" in "Using Message Filters to Enforce Email Policies," of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Reporting Incorrectly Classified Messages to Cisco IronPort Systems

Messages that appear to be incorrectly classified may be reported to Cisco IronPort for analysis. Each message is reviewed by a team of human analysts and used to enhance the accuracy and effectiveness of the product. Each message should be forwarded as an RFC 822 attachment to the following addresses:

- spam@access.ironport.com - for reporting missed spam
- ham@access.ironport.com - for reporting false-positives

Due to the volume of submissions, Cisco IronPort cannot provide individual feedback or results to customers.

For more information about reporting incorrectly classified messages, please see the Cisco IronPort Knowledge base or contact your Cisco IronPort Support provider.

Testing Cisco IronPort Anti-Spam

Step 1 Enable Cisco IronPort Anti-Spam on a mail policy (as above).

Step 2 Send a test email that includes the following header to a user in that mail policy: `X-Advertisement: spam`

For testing purposes, Cisco IronPort Anti-Spam considers any message with an X-header formatted as `X-Advertisement: spam` to be spam. The test message you send with this header is flagged by Cisco IronPort Anti-Spam, and you can confirm that the actions you configured for the mail policy (Configuring Per-Recipient Policies for Anti-Spam, page 9-12) are performed. You can use the `trace` command and include this header, or use a Telnet program to send SMTP commands to the appliance. See the "Testing and Troubleshooting" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide* and Appendix A, "Accessing the Appliance" for more information.

Note Examining a message's headers for specific headers added by Cisco IronPort Anti-Spam is another method to test the configuration of Cisco IronPort Anti-Spam on your appliance. See Headers Added by Cisco IronPort Anti-Spam and Intelligent Multi-Scan, page 9-16.

Evaluating Anti-Spam Efficacy

Cisco strongly recommends evaluating the product using a live mail stream directly from the Internet. This is because Cisco IronPort Anti-Spam and Cisco IronPort Intelligent Multi-Scan rules are added quickly to prevent active spam attacks and quickly expire once attacks have passed. Testing using old messages will therefore lead to inaccurate test results.

Using the `X-Advertisement: spam` header is the best method to test if your system configuration is correctly handling a message that would be considered spam if it were "live." Use the `trace` command (see Debugging Mail Flow Using Test Messages: Trace, page -446) or see the following example.

Common pitfalls to avoid while evaluating include:

- Evaluating using resent or forwarded mail or cut-and-pasted spam messages

  Mail lacking the proper headers, connecting IP, signatures, etc. will result in inaccurate scores.

- Testing "hard spam" only

  Removing the "easy spam" using SBRS, blacklists, message filters, etc. will result in a lower overall catch rate percentage.

- Resending spam caught by another anti-spam vendor
- Testing older messages

  CASE adds and removes rules rapidly based on current threats. Testing using an older collection of messages will significantly distort the results.

Example

Use SMTP commands to send a test message with the `X-advertisement: spam` header to an address to which you have access. Ensure that the mail policy is configured to receive messages for the test address (see Accepting Email for Local Domains or Specific Users on Public Listeners (RAT), page 5-50) and that the HAT will accept the test connection.

```
# telnet IP_address_of_IronPort_Appliance_with_IronPort_Anti-Spam port

220 hostname ESMTP

helo example.com

250 hostname

mail from: <test@example.com>

250 sender <test@example.com> ok

rcpt to: <test@address>

250 recipient <test@address> ok

data

354 go ahead

Subject: Spam Message Test

X-Advertisement: spam



spam test

.
```

```
250 Message MID accepted

221 hostname

quit
```

Then, check the mailbox of the test account and confirm that the test message was correctly delivered based upon the actions you configured for the mail policy.

For example:

- Was the subject line altered?
- Was your additional custom header added?
- Was the message delivered to an alternate address?
- Was the message dropped?

Incoming Relays

The Incoming Relays feature helps your Cisco IronPort appliance obtain the IP address of an external machine that is sending mail to the Cisco IronPort appliance via one or more mail exchange/transfer agents (MX or MTA), filtering servers, etc. at the edge of the network. In this type of configuration, the IP address of the external machine is not automatically known by the Cisco IronPort appliance. Instead, mail appears to originate from the local MX/MTA (the incoming relay) rather than from the external machine. Cisco IronPort Anti-Spam and Cisco IronPort Intelligent Multi-Scan depend on accurate IP addresses for external senders so it is vital for the Cisco IronPort appliance to have this information.

Note You should only enable this feature if you have a local MX/MTA relaying mail to your Cisco IronPort appliance.

Figure 9-12 shows a very basic example of an incoming relay. Mail from IP address 7.8.9.1 appears to come from IP address 10.2.3.4 because the local MX/MTA is relaying mail to the Cisco IronPort appliance.

Figure 9-12 Mail Relayed by MX/MTA — Simple

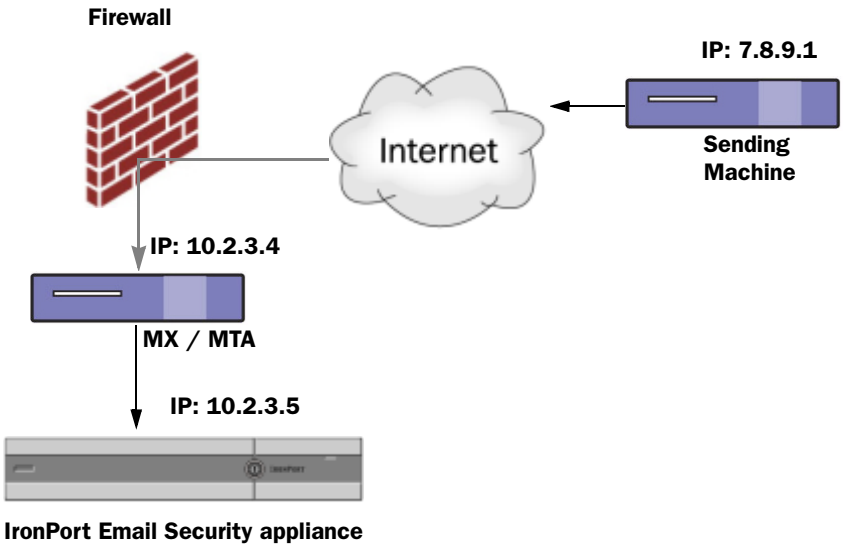

Figure 9-13 shows two other, slightly more complicated examples of how mail may be relayed inside the network and how mail may be processed by several servers within the network before it is passed to the Cisco IronPort appliance. In example A, mail from 7.8.9.1 passes through the firewall and is processed by an MX and an MTA before being delivered to the Cisco IronPort appliance. In example B, mail from 7.8.9.1 is sent to a load balancer or other type of traffic shaping appliance and is sent to any one of a range of MXs prior to being delivered to the Cisco IronPort appliance.

Figure 9-13 Mail Relayed by MX/MTA — Advanced

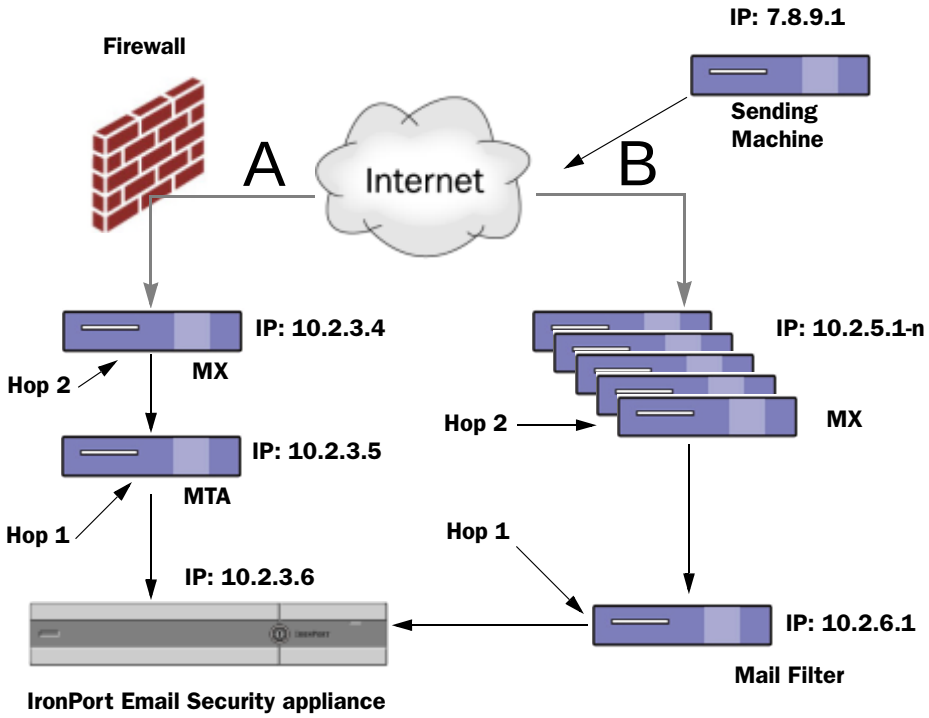

The Incoming Relays Feature: Overview

Occasionally, administrators need to run the Cisco IronPort appliance behind the mail exchange (MX) or mail transfer agent (MTA) at the edge of the network instead of receiving mail directly from the Internet. Unfortunately, when using this configuration the Cisco IronPort appliance is not receiving the mail directly from the Internet and so it does not have access to the last connecting IP address from the external network. Instead mail received is listed as being received from the local MX/MTA. It is critical for successful operation of the Cisco IronPort appliance that the connecting IP address be known so that SenderBase Reputation Service can be used in Cisco IronPort Intelligent Multi-Scan and Cisco IronPort Anti-Spam scanning.

The solution is to configure an incoming relay. When configuring an incoming relay, you specify the names and IP addresses of all of the internal MX/MTAs connecting to the Cisco IronPort appliance, as well as the header used to store the originating IP address. You can specify either an Internet Protocol version 4 (IPv4) or version 6 (IPv6) address for the internal MX/MTA. You have two options for specifying the header: a custom header or an existing received header.

Incoming Relays and Email Security Monitor

When using the Incoming Relay feature, data provided by the Email Security Monitor will contain data for both the external IP and the MX/MTA. For example, if an external machine (IP 7.8.9.1) sent 5 emails through the internal MX/MTA (IP 10.2.3.4), Mail Flow Summary will show 5 messages coming from IP 7.8.9.1 and 5 more coming from the internal relay MX/MTA (IP 10.2.3.5).

Incoming Relays and Filters

The Incoming Relays feature provides the various SenderBase Reputation Service related filter rules (`reputation, no-reputation`) with the correct SenderBase Reputation score.

Incoming Relays, HAT, SBRS, and Sender Groups

Please note that HAT policy groups do not currently use information from Incoming Relays. However, because the Incoming Relays feature does supply the SenderBase Reputation score, you can simulate HAT policy group functionality via message filters and the `$reputation` variable.

Incoming Relays and Reporting

When using Incoming Relays, the SenderBase Reputation score is not reported correctly in the Email Security Monitor reports. Also, sender groups may not be resolved correctly.

Incoming Relays and Message Tracking

When using Incoming Relays, the Message Tracking Details page displays the relay's IP address and the relay's SenderBase Reputation Score for a message instead of the IP address and reputation score of the sender.

Incoming Relays and Trace

Trace returns the Incoming Relay's SenderBase Reputation Score in its results instead of the reputation score for the source IP address.

Incoming Relays and Directory Harvest Attack Prevention

If a remote host attempts a directory harvest attack by sending messages to the MX or MTA serving as an incoming realy on your network, the appliance drops the connection from the incoming relay if the relay is assigned to a sender group with a mail flow policy with Directory Harvest Attack Prevention (DHAP) enabled. This prevents all messages from the relay, including legitimate messages, from reaching the Email Security applianc. The appliance does not have the opportunity to recognize the remote host as the attacker and the MX or MTA that's acting as the incoming relay continues to receive mail from the attacking host. To work around this issue and continue receiving messages from the incoming relay, add the relay to a sender group with a mail flow policy that has unlimited messages for DHAP.

IP Addresses

As a general rule, when specifying an IP address (of the machine connecting to the Cisco IronPort appliance — the incoming relay), be as specific as possible. That said, IP addresses can also be entered using standard CIDR format or an IP address range. For example, if you have several MTAs at the edge of your network receiving email, you might want to enter a range of IP addresses to include all of your MTAs, such as 10.2.3.1/8 or 10.2.3.1-10. You can use IPv4 or IPv6 addresses for the MTAs.

For IPv6 addresses, AsyncOS supports the following formats:

- `2620:101:2004:4202::0-2620:101:2004:4202::ff`
- `2620:101:2004:4202::`
- `2620:101:2004:4202::23`
- `2620:101:2004:4202::/64`

Message Headers and Incoming Relays

Custom Header

Use this method to specify a custom header. This is the recommended method. The machine connecting to the original sender needs to add this custom header. The value of the header is expected to be the IP address of the external sending machine. For example:

```
SenderIP: 7.8.9.1
```

```
X-CustomHeader: 7.8.9.1
```

When entering a header, you do not need to enter the trailing colon.

If your local MX/MTA can receive mail from a variable number of hops, inserting a custom header is the only way to enable the Incoming Relays feature. For example, in Figure 9-14 both path C and D lead to IP address 10.2.3.5; however, path C has two hops and path D has one. Because the number of hops can vary in this situation, you must use a custom header in order to have Incoming Relays configured correctly.

Figure 9-14 Mail Relayed by MX/MTA — Variable Number of Hops

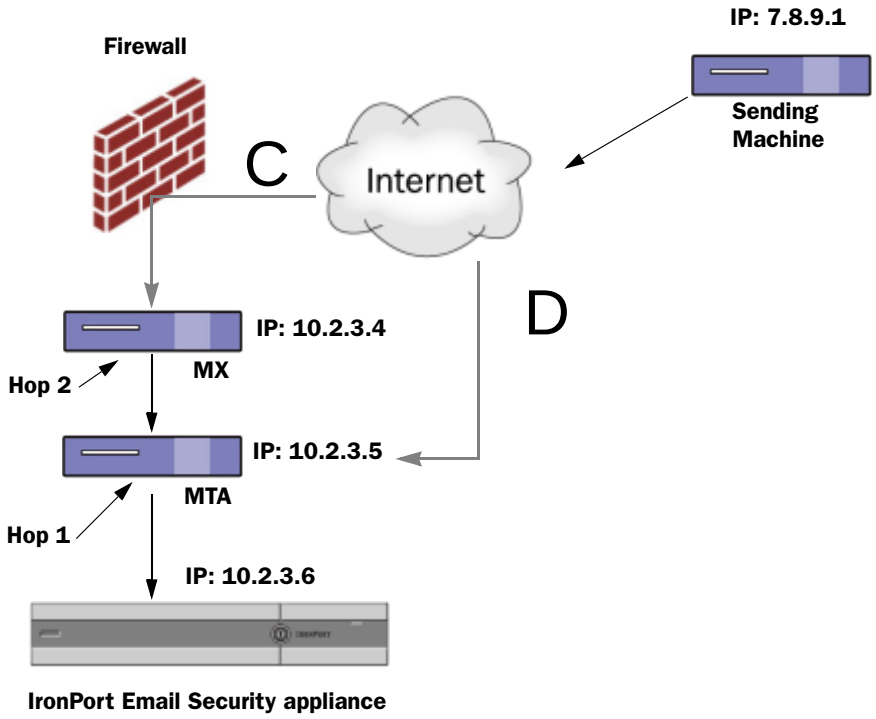

Received Header

If configuring the MX/MTAs to include a custom header containing the sending IP address is not an option, you can configure the incoming relays feature to attempt to determine the sending IP address by examining the "Received:" headers in the message. Using the "Received:" header will only work if the number of network "hops" will always be constant for an IP address. In other words, the machine at the first hop (10.2.3.5 in Figure 9-13) should always be the same number of hops away from the edge of your network. If incoming mail can take different paths (resulting in a different number of hops, as described in Figure 9-14) to the machine connecting to your Cisco IronPort appliance, you must use a custom header (see Custom Header, page 9-22).

Specify a parsing character or string and the number of network hops (or Received: headers) back to look. A hop is basically the message travelling from one machine to another (being received by the Cisco IronPort appliance does not count as a hop. See Determining Which Headers are Used, page 9-25 for more information). AsyncOS looks for the first IP address following the first occurrence of the parsing character or string in the Received: header corresponding to the number of specified hops. For example, if you specify two hops, the second Received: header, working backward from the Cisco IronPort appliance is parsed. If the parsing character is not found, or if there is not a valid IP address found, the Cisco IronPort appliance uses the real IP address of the connecting machine.

If you specify an opening square bracket ([ ) and two hops for the following example mail headers, the IP address of the external machine is 7.8.9.1. However, if you specify an closing parenthesis ( )) as the parsing character, a valid IP address will not be found. In this case, the Incoming Relays feature is treated as disabled, and the IP of the connecting machine is used (10.2.3.5).

In the example in Figure 9-13 the incoming relays are:

- Path A — 10.2.3.5 (with 2 hops when using received headers) and
- Path B — 10.2.6.1 (with 2 hops when using received headers)

Table 9-2 shows example email headers for a message as it moves through several hops on its way to the Cisco IronPort appliance as in Figure 9-13. This example shows extraneous headers (ignored by your Cisco IronPort appliance) which are present once the message has arrived in the recipient's inbox. The number of hops to specify would be two. Table 9-3 shows the headers for the same email message, without the extraneous headers

Table 9-2 A Series of Received: Headers (Path A Example 1)

1	Microsoft Mail Internet Headers Version 2.0 Received: from smemail.rand.org ([10.2.2.7]) by smmail5.customerdoamin.org with Microsoft SMTPSVC(5.0.2195.6713); Received: from ironport.customerdomain.org ([10.2.3.6]) by smemail.customerdoamin.org with Microsoft SMTPSVC(5.0.2195.6713);
2	Received: from mta.customerdomain.org ([10.2.3.5]) by ironport.customerdomain.org with ESMTP; 21 Sep 2005 13:46:07 -0700
3	Received: from mx.customerdomain.org (mx.customerdomain.org) [10.2.3.4]) by mta.customerdomain.org (8.12.11/8.12.11) with ESMTP id j8LKkWu1008155 for <joefoo@customerdomain.org>
4	Received: from sending-machine.spamham.com (sending-machine.spamham.com [7.8.9.1]) by mx.customerdomain.org (Postfix) with ESMTP id 4F3DA15AC22 for <joefoo@customerdomain.org>
5	Received: from linux1.thespammer.com (HELO linux1.thespammer.com) ([10.1.1.89]) by sending-machine.spamham.com with ESMTP; Received: from exchange1.thespammer.com ([10.1.1.111]) by linux1.thespammer.com with Microsoft SMTPSVC(6.0.3790.1830); Subject: Would like a bigger paycheck? Date: Wed, 21 Sep 2005 13:46:07 -0700 From: "A. Sender" <asend@otherdomain.com> To: <joefoo@customerdomain.org>

Notes for Table 9-2:

Step 1 The Cisco IronPort appliance ignores these headers.

Step 2 The Cisco IronPort appliance receives the message (not counted as a hop).

Step 3 First hop (and incoming relay).

Step 4 Second hop. This is the sending MTA. The IP address is 7.8.9.1.

Step 5 The Cisco IronPort appliance ignores these Microsoft Exchange headers.

Table 9-3 A Series of Received: Headers (Path A Example 2)

1	Received: from mta.customerdomain.org ([10.2.3.5]) by ironport.customerdomain.org with ESMTP; 21 Sep 2005 13:46:07 -0700

Table 9-3 A Series of Received: Headers (Path A Example 2) (Continued)

2	`Received: from mx.customerdomain.org (mx.customerdomain.org) [10.2.3.4]) by mta.customerdomain.org (8.12.11/8.12.11) with ESMTP id j8LKkWu1008155 for <joefoo@customerdomain.org>;`
3	`Received: from sending-machine.spamham.com (sending-machine.spamham.com [7.8.9.1]) by mx.customerdomain.org (Postfix) with ESMTP id 4F3DA15AC22 for <joefoo@customerdomain.org>;`

Figure 9-15 shows the incoming relay for path A (above) as configured in the Add Relay page in the GUI:

Figure 9-15 A Configured Incoming Relay

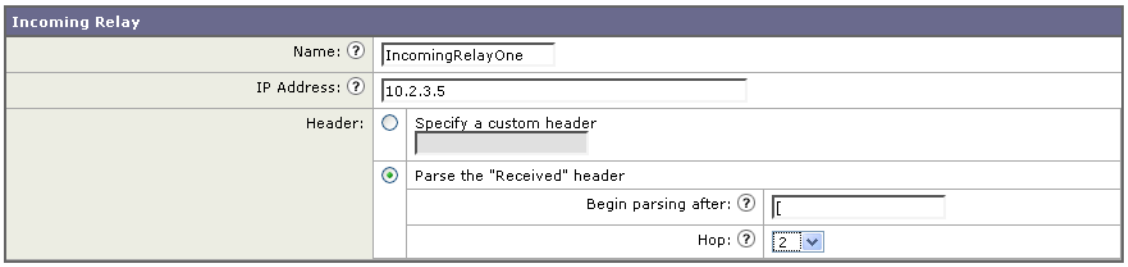

Determining Which Headers are Used

Your Cisco IronPort appliance will only examine the headers that were present when the message was received. So, additional headers added locally (such as Microsoft Exchange headers, etc.) or when the message is received by the Cisco IronPort appliance are not processed. One way to help determine which headers are used is to configure AsyncOS logging to include received headers via the `logheaders` subcommand of the `logconfig` CLI command:

```
mail3.example.com> logconfig


Currently configured logs:


[ ... list of configured logs ... ]


Choose the operation you want to perform:

- NEW - Create a new log.

- EDIT - Modify a log subscription.

- DELETE - Remove a log subscription.

- SETUP - General settings.
```

```
- LOGHEADERS - Configure headers to log.

- HOSTKEYCONFIG - Configure SSH host keys.

- CLUSTERSET - Set how logs are configured in a cluster.

- CLUSTERSHOW - Display how logs are configured in a cluster.

[]> logheaders


Please enter the list of headers you wish to record in the log files.

Separate multiple headers with commas.

[]> Received
```

Configuring the Incoming Relays Feature (GUI)

The Incoming Relays page is available via the Network tab.

Enabling the Incoming Relays Feature Globally

Step 1 Click the Incoming Relays link on the Network Tab. The Incoming Relays page is displayed:

Figure 9-16 Incoming Relays Page

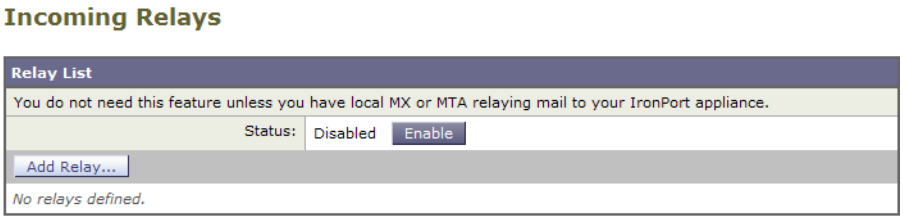

Step 2 Click **Enable** to enable Incoming Relays. (If the Incoming Relays feature is enabled, you can disable it by clicking **Disable.**)

Step 3 Commit your changes.

Incoming Relays and Mail Logs

The following example shows a typical log entry containing Incoming Relay information:

```
Wed Aug 17 11:20:41 2005 Info: MID 58298 IncomingRelay(myrelay): Header Received found,
IP 192.168.230.120 being used
```

Adding a Relay

Step 1 Click the **Add Relay** button on the Incoming Relays page. The Add Relay page is displayed:

Figure 9-17 Add Relay Page

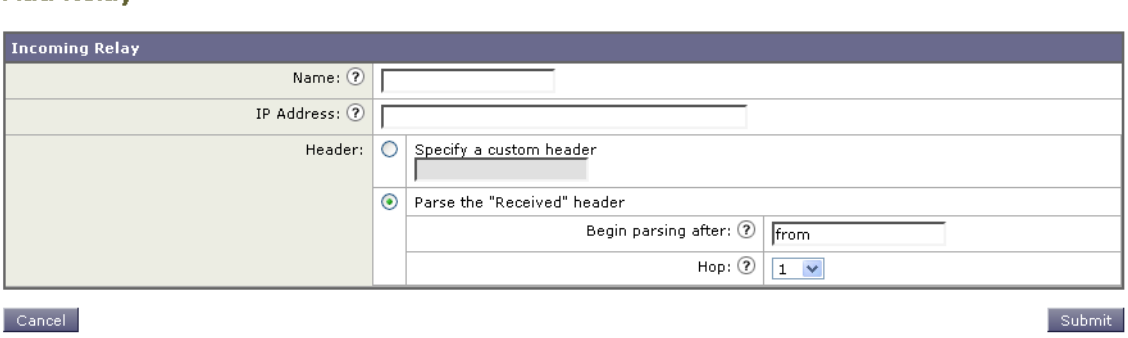

Step 2 Enter a name for the relay.

Step 3 Enter an IP address for the relay. For more information about valid IP address entries, including IPv6 address formatting, see IP Addresses, page 9-22.

Step 4 Select a header type (Custom or Received). For more information about custom headers, see Custom Header, page 9-22. When entering a header, you do not need to enter the trailing colon.

- For custom headers, enter the header name.
- For Received: headers, enter the character or string after which the IP address will appear. Enter the number for the "hop" to check for the IP address. For more information, see Received Header, page 9-23.

Step 5 Commit your changes.

Editing a Relay

Step 1 Click on the relay's name in the Incoming Relay page. The Edit Relay page is displayed.

Step 2 Make changes to the relay.

Step 3 Commit your changes.

Deleting a Relay

Step 1 Click on the trash can icon in the corresponding row for the relay you want to delete. You are prompted to confirm the deletion.

Step 2 Click **Delete**.

Step 3 Commit your changes.

Incoming Relays and Logging

In the following log example, the SenderBase Reputation score for the sender is reported initially on line 1. Later, once the Incoming Relay is processed, the correct SenderBase Reputation score is reported on line 5.

1	Fri Apr 28 17:07:29 2006 Info: ICID 210158 ACCEPT SG UNKNOWNLIST match nx.domain SBRS rfc1918
2	Fri Apr 28 17:07:29 2006 Info: Start MID 201434 ICID 210158
3	Fri Apr 28 17:07:29 2006 Info: MID 201434 ICID 210158 From: <joe@sender.com>
4	Fri Apr 28 17:07:29 2006 Info: MID 201434 ICID 210158 RID 0 To: <mary@example.com>
5	Fri Apr 28 17:07:29 2006 Info: MID 201434 IncomingRelay(senderdotcom): Header Received found, IP 192.192.108.1 being used, **SBRS 6.8**
6	Fri Apr 28 17:07:29 2006 Info: MID 201434 Message-ID '<7.0.1.0.2.20060428170643.0451be40@sender.com>'
7	Fri Apr 28 17:07:29 2006 Info: MID 201434 Subject 'That report...'
8	Fri Apr 28 17:07:29 2006 Info: MID 201434 ready 2367 bytes from <joe@sender.com>
9	Fri Apr 28 17:07:29 2006 Info: MID 201434 matched all recipients for per-recipient policy DEFAULT in the inbound table
10	Fri Apr 28 17:07:34 2006 Info: ICID 210158 close
11	Fri Apr 28 17:07:35 2006 Info: MID 201434 using engine: CASE spam negative
12	Fri Apr 28 17:07:35 2006 Info: MID 201434 antivirus negative
13	Fri Apr 28 17:07:35 2006 Info: MID 201434 queued for delivery

CHAPTER 10

Outbreak Filters

Low-volume, targeted email attacks such as phishing messages, scams, and malware links are on the rise while viruses spread through attachments are on the decline. The messages used for these non-viral attacks are complex and evolving; they are professional-looking messages that use social engineering tricks, including using the recipient's information, in an attempt to trick the recipient into clicking custom URLs that point to phishing and malware websites. These URLs can be unique for each recipient or a small group of recipients and these websites are online only for a short period of time and are unknown to web security services. All of these factors make these small scale, non-viral outbreaks more difficult to detect than widespread virus outbreaks and spam campaigns. Cisco IronPort's Outbreak Filters feature protects your users from this growing trend of targeted attacks in addition to new virus outbreaks.

Outbreak Filters Overview

Messages designed to steal sensitive information from users or deliver malware to their computers continue to evolve and can slip by traditional anti-virus and anti-spam scanning software. Outbreak Filters act proactively to provide a critical first layer of defense against these new outbreaks. By detecting new outbreaks in real-time and dynamically responding to prevent suspicious traffic from entering the network, Cisco IronPort's Outbreak Filters feature offers protection until new anti-virus and anti-spam updates are deployed. The Outbreak Filters use Cisco IronPort's outbreak detection technology and intelligent quarantine system to protect your users.

The Outbreak Filters feature protects your users and your network by gathering information about outbreaks as they occur and using this data to prevent the spread of these outbreaks to your users. Outbreak Filters compares incoming messages with published Outbreak Rules from Cisco Security Intelligence Operations (SIO) to determine if the message is a part of a large-scale virus outbreak or a smaller, non-viral attack. AsyncOS assigns messages that match the Outbreak Rules a threat level that indicates the severity of the message's threat and compares that threat level to the quarantine and message modfication thresholds you set for your mail policy. Messages that meet or exceed one of those thresholds are quarantined or modified to protect the recipient.

The process of outbreak detection and filtering begins with SenderBase, part of SIO. SenderBase is the world's largest email and web traffic monitoring system and has a view into approximately 25% of the world's email traffic. Cisco IronPort uses historical SenderBase data to create a statistical view of normal global traffic patterns. Outbreak Filters depends on the set of rules developed from this data to determine the threat levels of incoming messages.

Outbreak Filters has significant enhancements in features and usability. At a high level the enhancements include, but are not limited to:

- The increased threat types detected by Cisco Security Intelligence Operations (SIO) and used to create Outbreak Rules to detect non-viral attacks, such as phishing scams and malware distribution, in addition to virus outbreaks.
- CASE (Context Adaptive Scanning Engine) scanning that scans for URLs to detect non-viral threats, in addition to combining content analysis from Adaptive Rules and Outbreak Rules from SIO to detect outbreaks.
- Dynamic Quarantine, which re-evaluates messages periodically and auto-releases them from the quarantine based on Outbreak Rule updates.
- URL rewriting to redirect traffic to potentially harmful websites through the Cisco web security proxy, which either warns users that the website they are attempting to access may be malicious or blocks the website completely.

These feature enhancements are designed to increase the system's capture rate for outbreaks, provide enhanced visibility into an outbreak, and protect your users' computers and sensitive information.

Your Cisco IronPort appliance ships with a 30-day evaluation license for the Outbreak Filters feature.

Threat Categories

The Outbreak Filters feature provides protection from two categories of message-based outbreaks: *virus outbreaks*, which are messages with never-before-seen viruses in their attachments, and *non-viral threats*, which includes phishing attempts, scams, and malware distribution through links to an external website.

By default, the Outbreak Filters feature scans your incoming and outgoing messages for possible viruses during an outbreak. You can enable scanning for non-viral threats in addition to virus outbreaks if you enable anti-spam scanning on the appliance.

Note Your appliance needs a feature key for Cisco IronPort Anti-Spam or Cisco IronPort Intelligent Multi-Scan in order for Outbreak Filters to scan for non-viral threats.

Virus Outbreaks

The Outbreak Filters feature provides you with a head start when battling virus outbreaks. An outbreak occurs when messages with attachments containing never-before-seen viruses or variants of existing viruses spread quickly through private networks and the Internet. As these new viruses or variants hit the Internet, the most critical period is the window of time between when the virus is released and when the anti-virus vendors release an updated virus definition. Having advanced notice — even a few hours — is vital to curbing the spread of the malware or virus. During that vulnerability window, the newly-found virus can propagate globally, bringing email infrastructure to a halt.

Phishing, Malware Distribution, and Other Non-Viral Threats

Messages containing non-viral threats are designed to look like a message from a legitimate sources and often sent out to a small number of recipients. These messages may have one or more of the following characteristics in order to appear trustworthy:

- The recipient's contact information.
- HTML content designed to mimic emails from legitimate sources, such as social networks and online retailers.
- URLs pointing to websites that have new IP addresses and are online only for a short time, which means that email and web security services do not have enough information on the website to determine if it is malicious.
- URLs pointing to URL shortening services.

All of these characteristics make these messages more difficult to detect as spam. The Outbreak Filters feature provides a multi-layer defense from these non-viral threats to prevent your users from downloading malware or providing personal information to suspicious new websites.

If CASE discovers URLs in the message, it compares the message to existing Outbreak Rules to determine if the message is part of a small-scale non-viral outbreak and then assigns a threat level. Depending on the threat level, the Email Security appliance delays delivery to the recipient until more threat data can be gathered and rewrites the URLs in the message to redirect the recipient to the Cisco web security proxy if they attempt to access the website. The proxy displays a splash page warning the user that the website may contain malware.

Outbreak Filters - Multi-Layered Targeted Protection

The Outbreak Filters feature uses three tactics to protect your users from outbreaks:

- **Delay.** The Outbreak Filters feature delays messages that may be part of a virus outbreak or non-viral attack by quarantining the message. While quarantined, CASE receives updated Outbreak Rules and rescans the message to confirm whether any of them is part of an attack. CASE determines the rescan period based on the message's threat level. See Delaying Messages, page 10-4 for more information.
- **Redirect.** Based on the threat level, Outbreak Filters rewrites the URLs in non-viral attack messages to redirect the recipient through the Cisco web security proxy if they attempt to access any of the linked websites. The proxy displays a splash screen that warns the user that the website may contain malware, if the website is still operational, or displays an error message if the website has been taken offline. See Redirecting URLs, page 10-5 for more information on redirecting URLs.
- **Modify.** In addition to rewriting URLs in non-viral threat messages, Outbreak Filters can modify a message's subject and add a disclaimer above the message body to warn users about the message's content. See Modifying Messages, page 10-6 for more information.

Cisco Security Intelligence Operations

Cisco Security Intelligence Operations (SIO) is a security ecosystem that connects global threat information, reputation-based services, and sophisticated analysis to Cisco security appliances to provide stronger protection with faster response times.

SIO consists of three components:

- SenderBase. The world's largest threat monitoring network and vulnerability database.
- Threat Operations Center (TOC). A global team of security analysts and automated systems that extract actionable intelligence gathered by SenderBase.
- Dynamic Update. Real-time updates automatically delivered to Cisco IronPort appliances as outbreaks occur.

SIO compares real-time data from the global SenderBase network to common traffic patterns to identify anomalies that are proven predictors of an outbreak. TOC reviews the data and issues a threat level of the possible outbreak. Cisco IronPort Email Security appliances download updated threat levels and Outbreak Rules and use them to scan incoming and outgoing messages, as well as messages already in the Outbreak quarantine.

Information about current virus outbreaks can be found on SenderBase's website here:

`http://www.senderbase.org/`

The SIO website provides a list of current non-viral threats, including spam, phishing, and malware distribution attempts:

`http://tools.cisco.com/security/center/home.x`

Context Adaptive Scanning Engine

Outbreak Filters are powered by Cisco IronPort's unique Context Adaptive Scanning Engine (CASE). CASE leverages over 100,000 adaptive message attributes tuned automatically and on a regular basis, based on real-time analysis of messaging threats.

For virus outbreaks, CASE analyzes the message content, context and structure to accurately determine likely Adaptive Rule triggers. CASE combines Adaptive Rules and the real-time Outbreak Rules published by SIO to evaluate every message and assign a unique threat level.

To detect non-viral threats, CASE scans messages for URLs and uses Outbreak Rules from SIO to evaluate a message's threat level if one or more URLs are found.

Based on the message's threat level, CASE recommends a period of time to quarantine the message to prevent an outbreak. CASE also determines the rescan intervals so it can reevaluate the message based on updated Outbreak Rules from SIO. The higher the threat level, the more often it rescans the message while it is quarantined.

CASE also rescans messages when they're released from the quarantine. A message can be quarantined again if CASE determines that it is spam or contains a virus upon rescan.

For more information about CASE, see Cisco IronPort Anti-Spam and CASE: an Overview, page 9-4.

Delaying Messages

The period between when an outbreak or email attack occurs and when software vendors release updated rules is when your network and your users are the most vulnerable. A modern virus can propagate globally and a malicious website can deliver malware or collect your users' sensitive information during this period. Outbreak Filters protects your users and network by quarantining suspect messages for a limited period of time, giving Cisco and other vendors time to investigate the new outbreak.

When a virus outbreak occurs, suspicious messages with attachments are quarantined until updated Outbreak Rules and new anti-virus signatures prove the email's attachment is clean or a virus.

Small scale, non-viral threats contain URLs to malicious websites that may be online for a short period of time in order to evade detection by web security services or through URL shortening services in order to circumvent web security by putting a trustworthy website in the middle. By quarantining messages containing URLs that meet your threat level threshold, not only does CASE have the opportunity to reevaluate the message's content based on updated Outbreak Rules from SIO, but the messages can remain in the quarantine long enough that the linked website may go offline or be blocked by a web security solution.

See Dynamic Quarantine, page 10-9 more information on how Outbreak Filters quarantine suspicious messages.

Redirecting URLs

When CASE scans a message at the Outbreak Filters stage, it searches for URLs in the message body in addition to other suspicious content. CASE uses published Outbreak Rules to evaluate whether the message is a threat and then scores the message with the appropriate threat level. Depending on the threat level, Outbreak Filters protects the recipient by rewriting all the URLs to redirect the recipient to the Cisco web security proxy, except for URLs pointing to bypassed domains, and delaying the delivery of the message in order for TOC to learn more about the website if it appears to be part of a larger outbreak. See URL Rewriting and Bypassing Domains, page 10-16 for more information on bypassing URLs for trusted domains.

After the Email Security appliance releases and delivers the message, any attempt by the recipient to access the website is redirected through the Cisco web security proxy. This is an external proxy hosted by Cisco that displays a splash screen that warns the user that the website may be dangerous, if the website is still operational. If the website has been taken offline, the splash screen displays an error message.

If the recipient decides to click the message's URLs, the Cisco web security proxy displays a splash screen in the user's web browser to warn the user about the content of the message. Figure 10-1 shows an example of the splash screen warning. The recipient can either click **Ignore this warning** to continue on to the website or **Exit** to leave and safely close the browser window.

Figure 10-1 Cisco Security Splash Screen Warning

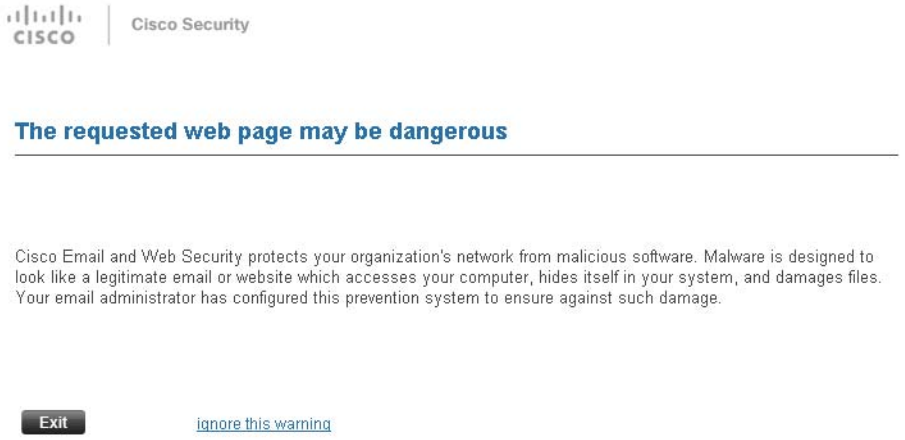

The only way to access the Cisco web security proxy is through a rewritten URL in a message. You cannot access the proxy by typing a URL in your web browser.

Modifying Messages

The Outbreak Filters feature modifies the message body of a non-viral threat message not only to rewrite the URLs but to alert the user that the message is a suspected threat. The Outbreak Filters feature can modify the subject header and add a disclaimer about the message's content above the message body. See Message Modification, page 10-16 for more information.

The threat disclaimer is created using the Disclaimer template through the Mail Policies > Text Resources page. See Managing Text Resources (GUI), page 14-13 for more information.

Types of Rules: Adaptive and Outbreak

Two types of rules are used by Outbreak Filters to detect potential outbreaks: Adaptive and Outbreak. The Outbreak Filters feature uses these two rule sets to provide the highest efficacy and the most focused set of criteria for threat detection to ensure that filters can be laser focused on a particular outbreak. The Outbreak Filters rules and actions are visible to the administrator, not hidden away behind the scenes, providing instant access to quarantined messages and the reason why they were quarantined.

Outbreak Rules

Outbreak Rules are generated by the Cisco IronPort Threat Operations Center (TOC), which is a part of the Cisco Security Intelligence Operations, and focus on the message as a whole, rather than just attachment filetypes. Outbreak Rules use SenderBase data (real time and historical traffic data) and any combination of message parameters such as attachment file type, file name keywords, or anti-virus engine update to recognize and prevent outbreaks in real time. Outbreak Rules are given a unique ID used to refer to the rule in various places in the GUI (such as the Outbreak quarantine).

Real-time data from the global SenderBase network is then compared to this baseline, identifying anomalies that are proven predictors of an outbreak. The TOC reviews the data and issues a threat indicator or Threat Level. The Threat Level is a numeric value between 0 (no threat) and 5 (extremely risky), and measures the likelihood that a message is a threat for which no other gateway defense is widely deployed by Cisco IronPort customers (for more information, see Threat Levels, page 10-7). Threat Levels are published as Outbreak Rules by the TOC.

Some example characteristics that can be combined in Outbreak Rules include:

- File Type, File Type & Size, File Type & File Name Keyword, etc.
- File Name Keyword & File Size
- File Name Keyword
- Message URL
- File Name & Sophos IDE

Adaptive Rules

Adaptive Rules are a set of rules within CASE that accurately compare message attributes to attributes of known virus outbreak messages. These rules have been created after studying known threat messages and known good messages within an extensive Cisco IronPort virus corpus. Adaptive Rules are updated often as the corpus is evaluated. They complement existing Outbreak Rules to detect outbreak messages at all times. While Outbreak Rules take effect when a possible outbreak is occurring, Adaptive Rules

(once enabled) are “always on,” catching outbreak messages locally before the full anomaly has formed on a global basis. Additionally, Adaptive Rules continuously respond to small and subtle changes in email traffic and structure, providing updated protection to customers.

Outbreaks

A Outbreak Filter rule is basically a Threat Level (e.g. 4) associated with a set of characteristics for an email message and attachment — things such as file size, file type, file name, message content, and so on. For example, assume the Cisco IronPort SIO notices an increase in the occurrences of a suspicious email message carrying a .exe attachment that is 143 kilobytes in size, and whose file name includes a specific keyword (“hello” for example). An Outbreak Rule is published increasing the Threat Level for messages matching this criteria. Your Cisco IronPort appliance checks for and downloads newly published Outbreak and Adaptive Rules every 5 minutes by default (see Updating Outbreak Filter Rules, page 10-13). Adaptive Rules are updated less frequently than Outbreak Rules. On the Cisco IronPort appliance, you set a threshold for quarantining suspicous messages. If the Threat Level for a message equals or exceeds the quarantine threshold, the message is sent to the *Outbreak* quarantine area. You can also set up a threshold for modifying non-viral threat messages to rewrite any URLs found in suspicious messages or add a notification at the top of message body.

Threat Levels

Table 10-1 on page 10-7 provides a basic set of guidelines or definitions for each of the various levels.

Table 10-1 Threat Level Definitions

Level	Risk	Meaning
0	None	There is no risk that the message is a threat.
1	Low	The risk that the message is a threat is low.
2	Low/Medium	The risk that the message is a threat is low to medium. It is a “suspected” threat.
3	Medium	Either the message is part of a confirmed outbreak or there is a medium to large risk of its content being a threat.
4	High	Either the message is confirmed to be part of a large scale outbreak or its content is very dangerous.
5	Extreme	The message’s content is confirmed to part of an outbreak that is either extremely large scale or large scale and extremely dangerous.

For more information about threat levels and outbreak rules, see Outbreak Filters Rules, page 10-13.

Guidelines for Setting Your Quarantine Threat Level Threshold

The quarantine threat level threshold allows administrators to be more or less aggressive in quarantining suspicious messages. A low setting (1 or 2) is more aggressive and will quarantine more messages; conversely, a higher score (4 or 5) is less aggressive and will only quarantine messages with an extremely high likelihood of being malicious.

The same threshold applies to both virus outbreaks and non-virus threats, but you can specify different quarantine retention times for virus attacks and other threats. See Dynamic Quarantine, page 10-9 for more information.

Cisco recommends the default value of 3.

Containers: Specific and Always Rules

Container files are files, such as zipped (.zip) archives, that contain other files. The TOC can publish rules that deal with specific files within archive files.

For example, if a virus outbreak is identified by TOC to consist of a .zip file containing a .exe, a specific Outbreak Rule is published that sets a threat level for .exe files within .zip files (.zip(exe)), but does not set a specific threat level for any other file type contained within .zip files (e.g. .txt files). A second rule (.zip(*)) covers all other file types within that container file type. An Always rule for a container will always be used in a message's Threat Level calculation regardless of the types of files that are inside a container. An always rule will be published by the SIO if all such container types are known to be dangerous.

Table 10-2 Fallback Rules and Threat Level Scores

Outbreak Rule	Threat Level	Description
.zip(exe)	4	This rule sets a threat level of 4 for .exe files within .zip files.
.zip(doc)	0	This rule sets a threat level of 0 for .doc files within .zip files.
zip(*)	2	This rule sets a threat level of 2 for all .zip files, regardless of the types of files they contain.

How the Outbreak Filters Feature Works

Email messages pass through a series of steps, the “email pipeline,” when being processed by your Cisco IronPort appliance (for more information about the email pipeline, see Understanding the Email Pipeline, page 4-1). As the messages proceed through the email pipeline, they are run through the anti-spam and anti-virus scanning engines if they are enabled for that mail policy. Only messages that pass through those scans are scanned by the Outbreak Filters feature (see Message and Content Filters and the Email Pipeline, page 10-21 for more information about how the email pipeline can affect which messages are scanned by the Outbreak Filters feature). In other words, known spam or messages containing recognized viruses are not scanned by the Outbreak Filters feature because they will have already been removed from the mail stream — deleted, quarantined, etc. — based on your anti-spam and anti-virus settings. Messages that arrive at the Outbreak Filters feature have therefore been marked spam- and virus-free. Note that a message quarantined by Outbreak Filters may be marked as spam or containing a virus when it is released from the quarantine and rescanned by CASE, based on updated spam rules and virus definitions.

Message Scoring

When a new virus attack or non-viral threat is released into the wild, no anti-virus or anti-spam software is able to recognize the threat yet, so this is where the Outbreak Filters feature can be invaluable. Incoming messages are scanned and scored by CASE using the published Outbreak and Adaptive Rules (see Types of Rules: Adaptive and Outbreak, page 10-6). The message score corresponds with the

message's threat level. Based on which, if any, rules the message matches, CASE assigns the corresponding threat level. If there is no associated threat level (the message does not match any rules), then the message is assigned a threat level of 0.

Once that calculation has been completed, the Email Security appliance checks whether the threat level of that message meets or exceeds your quarantine or message modification threshold value and quarantines message or rewrites its URLs. It the threat level is below the thresholds, it will be passed along for further processing in the pipeline.

Additionally, CASE reevaluates existing quarantined messages against the latest rules to determine the latest threat level of a message. This ensures that only messages that have a threat level consistent with an outbreak message stay within the quarantine and messages that are no longer a threat flow out of the quarantine after an automatic reevaluation.

In the case of multiple scores for an outbreak message — one score from an Adaptive Rule (or the highest score if multiple Adaptive Rules apply), and another score from an Outbreak Rule (or the highest score if multiple Outbreak Rules apply) — intelligent algorithms are used to determine the final threat level.

Note It is possible to use the Outbreak Filters feature without having enabled anti-virus scanning on the Cisco IronPort appliance. The two security services are designed to complement each other, but will also work separately. That said, if you do not enable anti-virus scanning on your Cisco IronPort appliance, you will need to monitor your anti-virus vendor's updates and manually release or re-evaluate some messages in the Outbreak quarantine. When using Outbreak Filters without anti-virus scanning enabled, keep the following in mind:

- You should disable Adaptive Rules
- Messages will get quarantined by Outbreak Rules
- Messages will get released if the threat level is lowered or time expires

Downstream anti-virus vendors (desktops/groupware) may catch the message on release.

Note Anti-spam scanning needs to be enabled globally on an appliance in order for the Outbreak Filters feature to scan for non-viral threats.

Dynamic Quarantine

The Outbreak Filters feature's Outbreak quarantine is a temporary holding area used to store messages until they're confirmed to be threats or it's safe to deliver to users. (See Outbreak Lifecycle and Rules Publishing, page 10-10 for more information.) Quarantined messages can be released from the Outbreak quarantine in several ways. As new rules are downloaded, messages in the Outbreak quarantine are reevaluated based on a recommended rescan interval calculated by CASE. If the revised threat level of a message falls under the quarantine retention threshold, the message will automatically be released (regardless of the Outbreak quarantine's settings), thereby minimizing the time it spends in the quarantine. If new rules are published while messages are being re-evaluated, the rescan is restarted.

Please note that messages quarantined as virus attacks are not automatically released from the outbreak quarantine when new anti-virus signatures are available. New rules may or may not reference new anti-virus signatures; however, messages will not be released due to an anti-virus engine update unless an Outbreak Rule changes the threat level of the message to a score lower than your Threat Level Threshold.

Messages are also released from the Outbreak quarantine after CASE's recommended retention period has elapsed. CASE calculates the retention period based on the message's threat level. You can define separate maximum retention times for virus outbreaks and non-viral threats. If CASE's recommended retention time exceeds the maximum retention time for the threat type, the Email Security appliance releases messages when the maximum retention time elapses. For viral messages the default maximum quarantine period is 1 day. The default period for quarantining non-viral threats is 4 hours. You can manually release messages from the quarantine.

The Email Security appliance also releases messages when the quarantine is full and more messages are inserted (this is referred to as overflow). Overflow only occurs when the Outbreak quarantine is at 100% capacity, and a new message is added to the quarantine. At this point, messages are released in the following order of priority:

- Messages quarantined by Adaptive Rules (those scheduled to be released soonest are first)
- Messages quarantined by Outbreak Rules (those scheduled to be released soonest are first)

Overflow stops the moment the Outbreak quarantine is below 100% capacity. For more information about how quarantine overflow is handled, see the "Quarantines" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.

Messages released from the Outbreak quarantine are scanned by the anti-virus and anti-spam engines again if they're enabled for the mail policy. If it is now marked as a known virus or spam, then it will be subject to your mail policy settings (including a possible second quarantining in the Virus quarantine or Cisco IronPort Spam quarantine). For more information, see The Outbreak Filters Feature and the Outbreak Quarantine, page 10-17.

Thus it is important to note that in a message's lifetime, it may actually be quarantined twice — once due to the Outbreak Filters feature, and once when it is released from the Outbreak quarantine. A message will not be subject to a second quarantine if the verdicts from each scan (prior to Outbreak Filters, and when released from the Outbreak quarantine) match. Also note that the Outbreak Filters feature does not take any final actions on messages. The Outbreak Filters feature will either quarantine a message (for further processing) or move the message along to the next step in the pipeline.

Outbreak Lifecycle and Rules Publishing

Very early in a virus outbreak's lifecycle, broader rules are used to quarantine messages. As more information becomes available, increasingly focused rules are published, narrowing the definition of what is quarantined. As the new rules are published, messages that are no longer considered possible virus messages are released from quarantine (messages in the outbreak quarantine are rescanned as new rules are published).

Table 10-3 Example Rules for an Outbreak Lifecycle

Time	Rule Type	Rule Description	Action
T=0	Adaptive Rule (based on past outbreaks)	A consolidated rule set based on over 100K message attributes, which analyzes message content, context and structure	Messages are automatically quarantined if they match Adaptive Rules
T=5 min	Outbreak Rule	Quarantine messages containing .zip (exe) files	Quarantine all attachments that are .zips containing a .exe

Table 10-3 Example Rules for an Outbreak Lifecycle (Continued)

Time	Rule Type	Rule Description	Action
T=10 min	Outbreak Rule	Quarantine messages that have .zip (exe) files greater than 50 KB	Any message with .zip (exe) files that are less than 50 KB would be released from quarantine
T=20 min	Outbreak Rule	Quarantine messages that have .zip (exe) files between 50 to 55 KB, and have "Price" in the file name	Any message that does not match this criteria would be released from quarantine
T=12 hours	Outbreak Rule	Scan against new signature	All remaining messages are scanned against the latest anti-virus signature

Managing Outbreak Filters (GUI)

Log in to the Graphical User Interface (GUI), select Security Services in the menu, and click Outbreak Filters.

Figure 10-2 Outbreak Filters Main Page

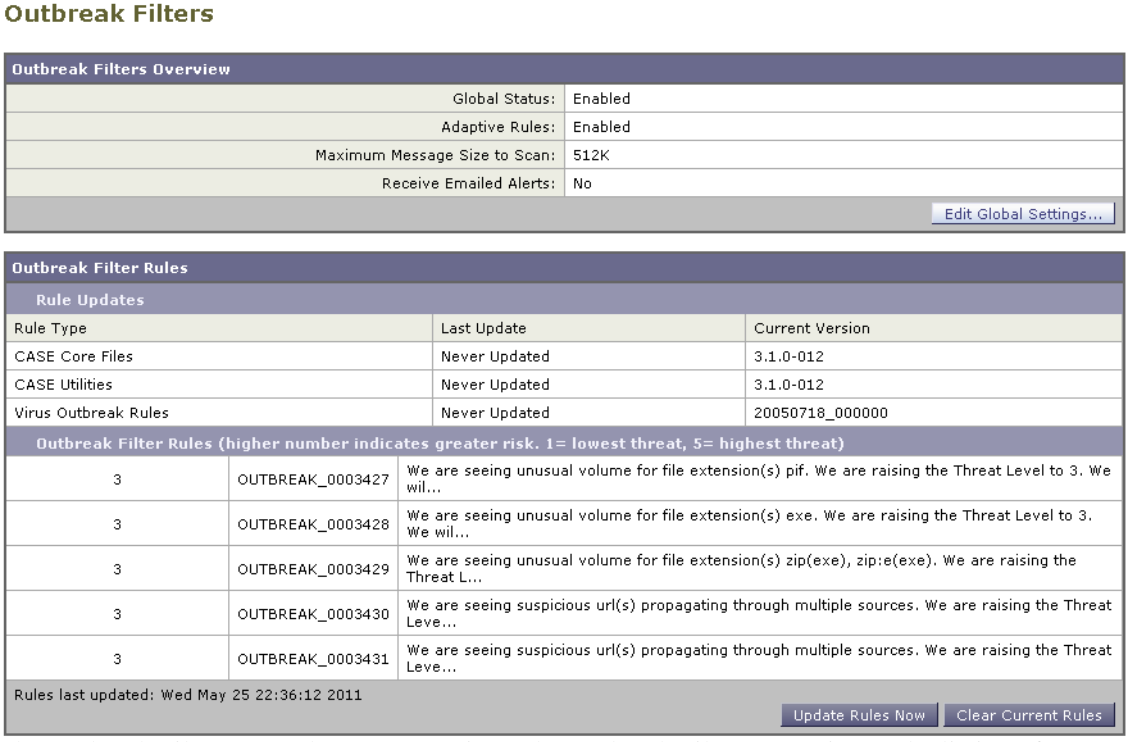

The Outbreak Filters page shows two sections: the Outbreak Filters Overview and a listing of current Outbreak Filter Rules (if any).

In Figure 10-2, Outbreak Filters are enabled, Adaptive Scanning is enabled, and the maximum message size is set to 512k. To change these settings, click **Edit Global Settings** For more information about editing Global Settings, see Configuring Outbreak Filters Global Settings, page 10-12.

The Outbreak Filter Rules section lists the time, date, and version of the latest update for various components (the rules engine as well as the rules themselves), as well as a listing of the current Outbreak Filter rules with threat level.

For more information about Outbreak Rules, see Outbreak Filters Rules, page 10-13.

Configuring Outbreak Filters Global Settings

To configure the Global Settings for Outbreak Filters, click **Edit Global Settings** The Outbreak Filters Global Settings page is displayed:

Figure 10-3 Outbreak Filters Global Settings Page

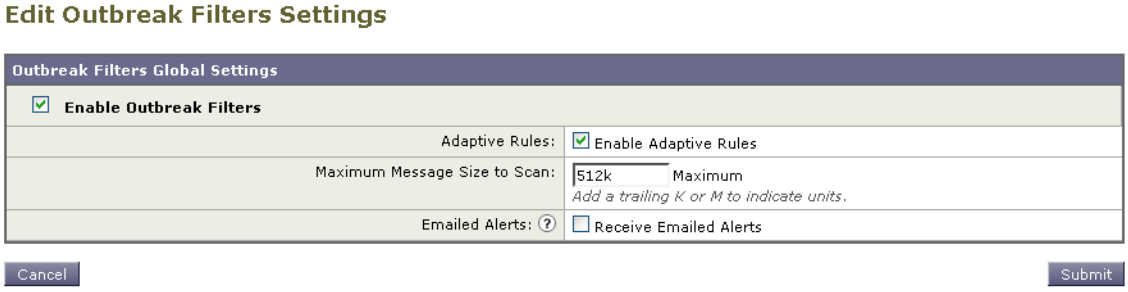

Use this page to:

- Enable Outbreak Filters globally.
- Enable Adaptive Rules scanning.
- Set a maximum size for files to scan (note that you are entering the size in *bytes*)
- Elect whether to enable alerts for the Outbreak Filter.

Note that alerts and Adaptive Rules are not enabled by default. This functionality is also available via the `outbreakconfig` CLI command (see the *Cisco IronPort AsyncOS CLI Reference Guide*). After you make your changes, submit and commit them.

Enabling the Outbreak Filters Feature

To enable the Outbreak Filters feature globally, check the box next to Enable Outbreak Filters on the Outbreak Filters Global Settings page, and click **Submit**. You must have agreed to the Outbreak Filters license agreement first.

Once enabled globally, the Outbreak Filters feature can then be enabled or disabled individually for each incoming and outgoing mail policy, including the default policies. For more information, see The Outbreak Filters Feature and Mail Policies, page 10-13.

The Outbreak Filters feature uses the Context Adaptive Scanning Engine (CASE) to detect viral threats, regardless of whether anti-spam scanning is enabled, but you do need to have Cisco IronPort Anti-Spam or Intelligent Multi-Scan enabled globally on the aplliance in order to scan for non-viral threats.

Note If you have not already agreed to the license during system setup (see Step 4: Security, page 3-21), you must click **Enable** on the Security Services > Outbreak Filters page, and then read and agree to the license.

Enabling Adaptive Rules

Adaptive Scanning enables the use of Adaptive Rules in Outbreak Filters. A set of factors or traits (file size, etc.) are used to determine the likelihood of a message being part of an outbreak when no virus signature or spam criteria relating to the message's content is available. To enable Adaptive Scanning, check the box next to Enable Adaptive Rules on the Outbreak Filters Global Settings page, and click **Submit**.

Enabling Alerts for Outbreak Filters

Check the box labeled "Emailed Alerts" to enable alerting for the Outbreak Filters feature. Enabling emailed alerts for Outbreak Filters merely enables the alerting engine to send alerts regarding Outbreak Filters. Specifying which alerts are sent and to which email addresses is configured via the Alerts page in the System Administration tab. For more information on configuring alerts for Outbreak Filters, see Alerts, SNMP Traps, and Outbreak Filters, page 10-20.

Outbreak Filters Rules

Outbreak Rules are published by the Cisco IronPort Security Intelligence Operations and your Cisco IronPort appliance checks for and downloads new outbreak rules every 5 minutes. You can change this update interval. See Editing Update Settings, page 15-11 for more information.

Managing Outbreak Filter Rules

Because the Outbreak Filters Rules are automatically downloaded for you, there really is no management needed on the part of the user.

However, if for some reason your Cisco IronPort appliance is not able to reach Cisco IronPort's update servers for new rules over a period of time, it is possible that your locally-cached scores are no longer valid, i.e., if a known viral attachment type now has an update in the anti-virus software and/or is no longer a threat. At this time, you may wish to no longer quarantine messages with these characteristics.

You can manually update the current outbreak rules by clicking **Update Rules Now**. This is identical to issuing the `outbreakupdate` command via the CLI (see the *Cisco IronPort AsyncOS CLI Reference Guide*).

Updating Outbreak Filter Rules

By default, your Cisco IronPort appliance will attempt to download new Outbreak Filters rules every 5 minutes. You can change this interval via the Security Services > Service Updates page. For more information, see Service Updates, page 15-10.

The Outbreak Filters Feature and Mail Policies

The Outbreak Filters feature has settings that can be set per mail policy. The Outbreak Filters feature can be enabled or disabled for each mail policy on the appliance. Specific file extensions and domains can be exempted from processing by the Outbreak Filters feature, per mail policy. This functionality is also available via the `policyconfig` CLI command (see the *Cisco IronPort AsyncOS CLI Reference Guide*).

Note Cisco IronPort Anti-Spam or Intelligent Multi-Scan scanning needs to be enabled globally on an appliance in order for the Outbreak Filters feature to scan for non-viral threats.

Figure 10-4 Mail Policy Listing

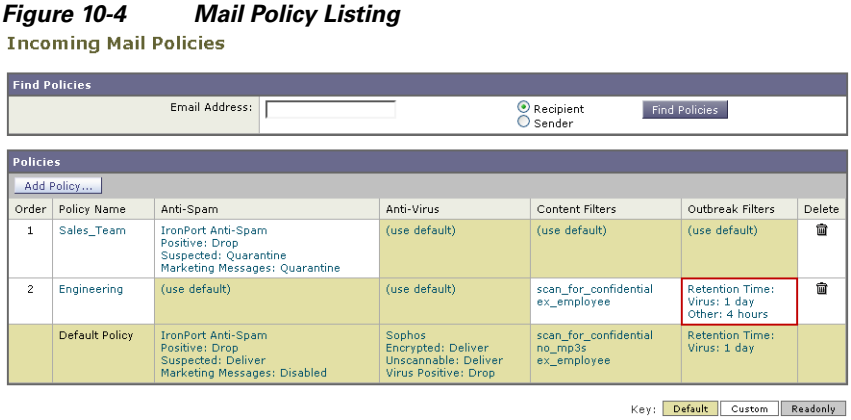

To modify the Outbreak Filters feature settings for a specific mail policy, click the link in the Outbreak Filters column of the policy to change. The Outbreak Filter Settings page is displayed.

Figure 10-5 Outbreak Filters Settings and Mail Policies

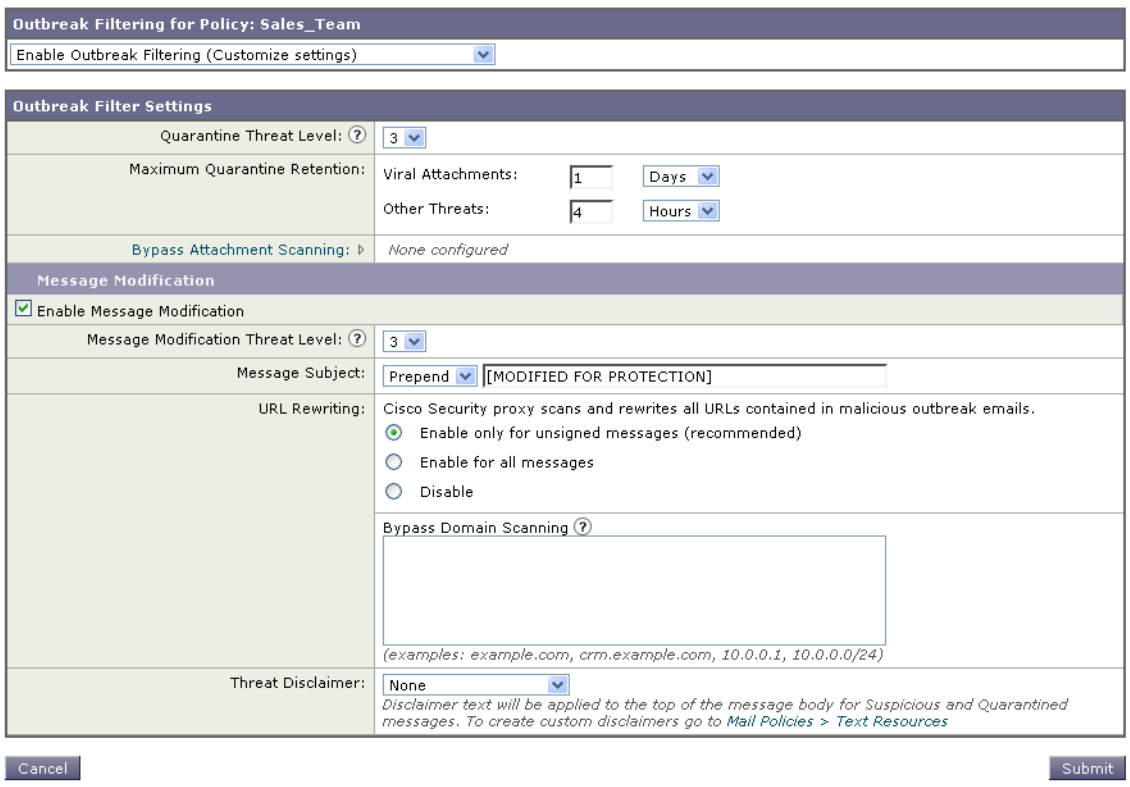

To enable and customize the Outbreak Filters feature for a particular mail policy, select **Enable Outbreak Filtering (Customize Settings)**.

You can configure the following Outbreak Filter settings for a mail policy:

- Quarantine threat level.
- Maximum quarantine retention time.
- File extension types for bypassing.
- Message modification threshold.
- Message subject.
- URL rewriting.
- Threat disclaimer.

Select **Enable Outbreak Filtering (Inherit Default mail policy settings)** to use the Outbreak Filters settings that are defined for the default mail policy. If the default mail policy has the Outbreak Filters feature enabled, all other mail policies use the same Outbreak Filter settings unless they are customized.

Once you have made your changes, commit your changes.

Setting a Quarantine Level Threshold

Select a Quarantine Threat Level threshold for outbreak threats from the list. A smaller number means that you will be quarantining more messages, while a larger number results in fewer messages quarantined. Cisco recommends the default value of 3.

For more information, see Guidelines for Setting Your Quarantine Threat Level Threshold, page 10-7.

Maximum Quarantine Retention

Specify the maximum amount of time in either hours or days that messages stay in the Outbreak Quarantine. You can specify different retention times for messages that may contain viral attachments and messages that may contain other threats, like phishing or malware links. You cannot quarantine non-viral threats unless you enable Message Modification for the policy.

CASE recommends a quarantine retention period when assigning the threat level to the message. The Email Security appliance keeps the message quarantined for the length of time that CASE recommends unless it exceeds the maximum quarantine retention time for its threat type.

Bypassing File Extension Types

You can modify a policy to bypass specific file types. Bypassed file extensions are not included when CASE calculates the threat level for the message; however, the attachments are still processed by the rest of the email security pipeline.

To bypass a file extension, click Bypass Attachment Scanning, select or type in a file extension, and click **Add Extension**. AsyncOS displays the extension type in the File Extensions to Bypass list.

To remove an extension from the list of bypassed extensions, click the trash can icon next to the extension in the File Extensions to Bypass list.

Bypassing File Extensions: Container File Types

When bypassing file extensions, files within container files (a .doc file within a .zip, for example) are bypassed if the extension is in the list of extensions to bypass. For example, if you add .doc to the list of extensions to bypass, all .doc files, even those within container files are bypassed.

Message Modification

Enable Message Modification if you want the appliance to scan messages for non-viral threats, such as phishing attempts or links to malware websites.

Based on the message's threat level, AsyncOS can modify the message to rewrite all of the URLs to redirect the recipient through the Cisco web security proxy if they attempt to open the website from the message. The appliance can also add a disclaimer to the message to alert the user that the message's content is suspicious or malicious.

You need to enable message modification in order to quarantine non-viral threat messages.

Message Modification Threat Level

Select a Message Modification Threat Level threshold from the list. This setting determines whether to modify a message based on the threat level returned by CASE. A smaller number means that you will be modifying more messages, while a larger number results in fewer messages being modified. Cisco recommends the default value of 3.

Message Subject

You can alter the text of the Subject header on non-viral threat messages containing modified links by prepending or appending certain text strings to notify users that the message has been modified for their protection.

Note White space is not ignored in the Message Subject field. Add spaces after (if prepending) or before (if appending) the text you enter in this field to separate your added text from the original subject of the message. For example, add the text `[MODIFIED FOR PROTECTION]` with a few trailing spaces if you are prepending.

Note The Message Subject field only accepts US-ASCII characters.

URL Rewriting and Bypassing Domains

If the message's threat level exceeds the message modification threshold, the Outbreak Filters feature rewrites all URLs in the message to redirect the user to the Cisco web security proxy's splash page if they click on any of them. (See Redirecting URLs, page 10-5 for more information.) If the message's threat level exceeds the quarantine threshold, the appliance also quarantines the message. If a small scale, non-viral outbreak is in progress, quarantining the message gives TOC time to analyze any suspect websites linked from possible outbreak messages and determine whether the websites are malicious. CASE uses updated Outbreak Rules from SIO to rescan the message to determine if it is part of the outbreak. After the retention period expires, the appliance releases the message from the quarantine.

AsyncOS rewrites all of the URLs inside a message except for the ones pointing to bypassed domains.

The following options are available for URL rewriting:

- **Enable only for unsigned messages**. This option allows AsyncOS to rewrite URLs in unsigned messages that meet or exceed the message modification threshold, but not signed messages. Cisco recommends using this setting for URL rewriting.

Note The Email Security appliance may rewrite URLs in a DomainKeys/DKIM-signed message and invalidate the message's signature if a server or appliance on your network other than the Email Security appliance is responsible for verifying the DomainKeys/DKIM signature.

- **Enable for all messages.** This option allows AsyncOS to rewrite URLs in all messages that meet or exceed the message modification threshold, including signed ones. If AsyncOS modifies a signed message, the signature becomes invalid.
- **Disable.** This option disables URL rewriting for Outbreak Filters.

You can modify a policy to exclude URLs to certain domains from modification. To bypass domains, enter the IPv4 address, IPv6 address, CIDR range, hostname, partial hostname or domain in the Bypass Domain Scanning field. Separate multiple entries using commas.

Threat Disclaimer

The Email Security appliance can append a disclaimer message above the heading of a suspicious message to warn the user of its content. This disclaimer can be in HTML or plain text, depending on the type of message.

Select the disclaimer text you want to use from the Threat Disclaimer list or click the Mail Policies > Text Resources link to create a new disclaimer using the Disclaimer Template. The Disclaimer Template includes variables for outbreak threat information. You can see a preview of the threat disclaimer by clicking Preview Disclaimer. For custom disclaimer messages, you can use variables to display the threat level, the type of threat, and a description of the threat in the message. For information on creating a disclaimer message, see Managing Text Resources (GUI), page 14-13.

The Outbreak Filters Feature and the Outbreak Quarantine

Messages quarantined by the Outbreak Filters feature are sent to the Outbreak quarantine. This quarantine functions like any other quarantine (for more information about working with quarantines, see the "Quarantines" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*) except that it has a "summary" view, useful for deleting or releasing all messages from the quarantine, based on the rule used to place the message in the quarantine (for Outbreak Rules, the Outbreak ID is shown, and for Adaptive Rules, a generic term is shown). For more information about the summary view, see Outbreak Quarantine and the Manage by Rule Summary View, page 10-19.

Figure 10-6 The Outbreak Quarantine

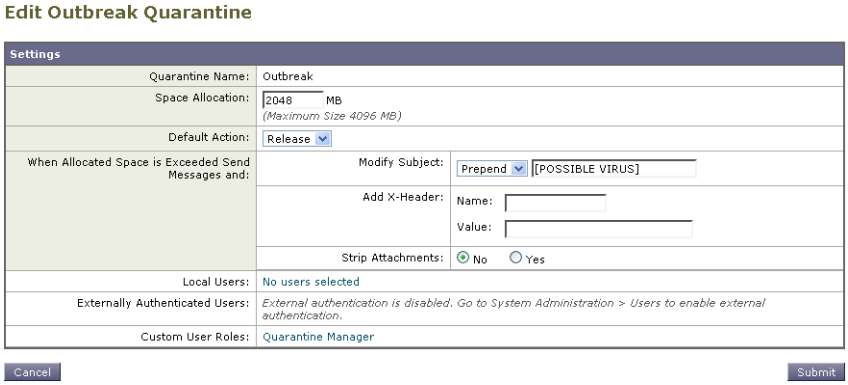

Monitoring the Outbreak Quarantine

Though a properly configured quarantine requires little if any monitoring, it is a good idea to keep an eye on the Outbreak Quarantine, especially during and after virus outbreaks when legitimate messages may be delayed.

If a legitimate message is quarantined, one of the following occurs depending on the settings for the Outbreak quarantine:

- If the quarantine's Default Action is set to Release, the message will be released when the retention time period expires or when the quarantine overflows. You can configure the Outbreak quarantine so that the following actions are performed on messages before they are released due to overflow: strip attachments, modify the subject, and add an X-Header. For more information about these actions, see the "Quarantines" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.
- If the quarantine's Default Action is set to Delete, the message will be deleted when the retention time period expires, or when the quarantine overflows.
- Overflow occurs when the quarantine is full and more messages are added. In this case the messages closest to their expiration date (not necessarily the oldest messages) are released first, until enough room is available for the new messages. You can configure the Outbreak quarantine so that the following actions are performed on messages before they are released due to overflow: strip attachments, modify the subject, add an X-Header.

Because quarantined messages are rescanned whenever new rules are published, it is very likely that messages in the Outbreak quarantine will be released prior to the expiration time.

Still, it can be important to monitor the Outbreak quarantine if the Default Action is set to Delete. Cisco recommends most users to not set the default action to Delete. For more information about releasing messages from the Outbreak quarantine, or changing the Default Action for the Outbreak Quarantine, see the "Quarantines" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.

Conversely, if you have messages in your Outbreak quarantine that you would like to keep in the quarantine longer while you wait for a new rule update, for example, you can delay the expiration of those messages. Keep in mind that increasing the retention time for messages can cause the size of the quarantine to grow.

Note If anti-virus scanning is disabled globally (not via a mail policy) while a message is in the Outbreak quarantine, the message is not anti-virus scanned when it leaves the quarantine, even if anti-virus scanning is re-enabled prior to the message leaving the quarantine.

Note You can use the Outbreak Filters feature without having enabled anti-virus scanning on the Cisco IronPort appliance. However, Outbreak Filters cannot scan for non-viral threats if anti-spam scanning is not enabled on the appliance.

Outbreak Quarantine and the Manage by Rule Summary View

You can view the contents of the Outbreak quarantine by clicking on the name of the quarantine in the listing on the Monitor menu in the GUI. The Outbreak quarantine has an additional view as well, the Outbreak Quarantine Manage by Rule Summary link.

Figure 10-7 The Outbreak Quarantine Manage by Rule Summary Link

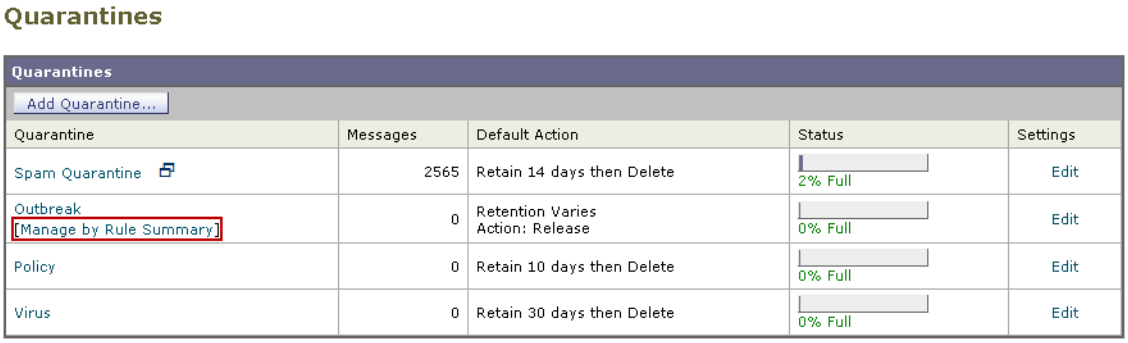

Using the Summary View to Perform Message Actions on Messages in the Outbreak Quarantine Based on Rule ID.

Click on the Manage by Rule Summary link to see a listing of the contents of the Outbreak quarantine, grouped by rule ID:

Figure 10-8 The Outbreak Quarantine Manage by Rule Summary View

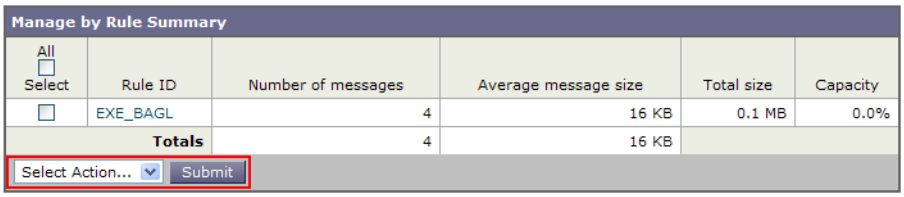

From this view, you can choose to release, delete, or delay the exit for all messages pertaining to a specific outbreak or adaptive rule, rather than selecting individual messages. You can also search through or sort the listing.

This functionality is also available via the `quarantineconfig -> outbreakmanage` CLI command. For more information, see the *Cisco IronPort AsyncOS CLI Reference Guide*.

Monitoring Outbreak Filters

The Cisco IronPort appliance includes several tools to monitor the performance and activity of the Outbreak Filters feature.

Outbreak Filters Report

The Outbreak Filters report to view the current status and configuration of Outbreak Filters on your Cisco IronPort appliance as well as information about recent outbreaks and messages quarantined due to Outbreak Filters. View this information on the Monitor > Outbreak Filters page. For more information, see the "Email Security Monitor" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.

Outbreak Filters Overview and Rules Listing

The overview and rules listing provide useful information about the current status of the Outbreak Filters feature. View this information via the Security Services > Outbreak Filters page.

Outbreak Quarantine

Use the outbreak quarantine to monitor how many messages are being flagged by your Outbreak Filters threat level threshold. Also available is a listing of quarantined messages by rule. View this information via the Monitor > Local Quarantines > Outbreak link and the Manage Rule by Summary link on the Monitor > Local Quarantines page. See the "Quarantines" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide* for more information.

Alerts, SNMP Traps, and Outbreak Filters

The Outbreak Filters feature supports two different types of notifications: regular AsyncOS alerts and SNMP traps.

SNMP traps are generated when a rule update fails. For more information about SNMP traps in AsyncOS, see the "Managing and Monitoring via the CLI" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.

AsyncOS has two types of alerts for the Outbreak Filter feature: size and rule

AsyncOS alerts are generated whenever the Outbreak quarantine's size goes above 5, 50, 75, and 95 of the maximum size. The alert generated for the 95% threshold has a severity of CRITICAL, while the remaining alert thresholds are WARNING. Alerts are generated when the threshold is crossed as the quarantine size increases. Alerts are not generated when thresholds are crossed as the quarantine size decreases. For more information about alerts, see Alerts, page 15-15.

AsyncOS also generates alerts when rules are published, the threshold changes, or when a problem occurs while updating rules or the CASE engine.

Troubleshooting The Outbreak Filters Feature

This section provides some basic troubleshooting tips for the Outbreak Filters feature.

Use the checkbox on the Manage Quarantine page for the Outbreak quarantine to notify Cisco of mis-classifications.

Multiple Attachments and Bypassed Filetypes

Bypassed file types are only excluded if a message's only attachment is of that type, or in the case of multiple attachments, if the other attachments do not yet have existing rules. Otherwise the message is scanned.

Message and Content Filters and the Email Pipeline

Message and content filters are applied to messages prior to scanning by Outbreak Filters. Filters can cause messages to skip or bypass the Outbreak Filters scanning.

CHAPTER 11

Data Loss Prevention

In the Information Age, your organization's data is one of its most prized possessions. Your organization spends a lot of money making data available to your employees, customers, and partners over email and the Web. This increased access poses challenges for information security professionals to figure out how to prevent the malicious or unintentional distribution of sensitive and proprietary information over the Internet.

Cisco provides the following methods to protect your organization's information and intellectual property and enforce compliancy with state and federal regulations using the Email Security appliance:

- **RSA Email DLP.** A solution local to the Email Security appliance that includes an integrated data loss prevention (DLP) scanning engine and DLP policy templates designed by RSA Security Inc. to identify and protect sensitive data.
- **RSA Enterprise Manager.** Users of RSA's Enterprise Manager can partner their Email Security appliances with the Enterprise Manager software and use RSA's DLP technologies to scan outgoing message. Whereas RSA Email DLP is local to an individual Email Security appliance, RSA Enterprise Manager allows you to manage multiple Email Security appliances on the same network from a centralized interface. Users of RSA's DLP Datacenter can use its fingerprinting detection method for scanning source code and documents in certain DLP policies. Enterprise Manager is a third-party software from RSA and cannot be purchased from Cisco.

Note This chapter describes how to configure the settings on the Email Security appliance to connect it to Enterprise Manager and provides an overview of how the appliance works as an Enterprise Manager partner device. For information on configuring the Enterprise Manager and its DLP policies, see RSA's documentation for Enterprise Manager, including the online help and the technical note *Managing Partner Device DLP with Enterprise Manager.*

Data Loss Prevention Overview

The Cisco IronPort Email Security appliance's Data Loss Prevention feature secures your organization's information and intellectual property and enforces regulatory and organizational compliance by preventing users from emailing sensitive data from your network. You define what kind of data your employees are not allowed to email by creating DLP policies that scan outgoing messages for any data that may violate laws or corporate policies.

This document refers to any message content that violates your DLP policies as *DLP violation* and the occurrence of message containing a violation as a *DLP incident*. When a DLP incident occurs, the appliance takes the appropriate actions with the message to secure the information, such as quarantining the message and sending a notification to someone in your organization responsible for data security.

The Email Security appliance has an integrated DLP scanning engine and a set of DLP policies created by RSA, which is referred to collectively in this documentation and on the appliance as RSA Email DLP. You can configure the Email Security appliance's outgoing mail policies to scan messages and attachments for DLP violations. RSA Email DLP includes over 100 DLP policy templates designed by RSA. See RSA Email DLP, page 11-8 for more information.

For users of RSA's Enterprise Manager, you can connect your Email Security appliances to Enterprise Manager as partner devices, allowing the appliances to use Enterprise Manager as a centralized management interface for multiple appliance on the network. Enterprise Manager provides a wider array of DLP technologies than RSA Email DLP does on the local Email Security appliance.

RSA Email DLP's policies are configured locally on the appliance while Enterprise Manager can manage the DLP policies for multiple Email Security appliances, including clustered appliances, and pushes those policies to the appliances for when the outgoing mail policies perform DLP scans.

If enabled, DLP scanning is performed in the appliance's "work queue" for outgoing mail immediately after the Outbreak Filters stage. See Message Splintering, page 6-4 for more information.

Data Loss Prevention Global Settings

To scan outgoing emails for sensitive data, you must first enable the Data Loss Prevention feature using the **Security Services > RSA Email DLP** page. You can choose whether to use RSA Enterprise Manager or RSA Email DLP for data loss prevention.

Select RSA Email DLP if you want to configure and manage your DLP policies on the local Email Security appliance. You can choose to either run the DLP Assessment Wizard to enable the most popular DLP policies on the appliance or manually configure DLP policies. To learn how to run the DLP Assessment Wizard, see Using the DLP Assessment Wizard, page 11-17. To learn how to manually configure DLP policies, see DLP Policy Manager, page 11-11.

After you enable RSA Email DLP, you can enable the policies on your outgoing mail policies using the Email Security Manager. For more information, see Configuring Per-Recipient Policies for DLP, page 11-31.

Select RSA Enterprise Manager if you want to use Enterprise Manager to configure and manage the DLP policies for your appliances. Enterprise Manager receives outgoing mail policy and message action definitions from the Email Security appliance and then pushes DLP policies to connected Email Security appliances. Administrators can also view DLP incidents and send commands to delete or release messages from quarantines using Enterprise Manager.

Both RSA Email DLP and RSA Enterprise Manager offer the option to log the content that violates your DLP policies, along with the surround content, which can then be viewed in the Message Tracking. This content may include sensitive data such as credit card numbers and social security numbers. Do not select this option if you don't want the appliance to log this information.

You can switch back to managing data loss prevention on the local appliance using RSA Email DLP whenever you want.

Enabling RSA Email DLP

Note If you want to use the DLP Assessment Wizard to configure the appliance's DLP policies, see Using the DLP Assessment Wizard, page 11-17.

Step 1 Select Security Services > RSA Email DLP.

Step 2 Click **Enable**.

Step 3 The license agreement page is displayed.

Note If you do not accept the license agreement, RSA Email DLP is not enabled on the appliance.

Step 4 Scroll to the bottom of the page and click **Accept** to accept the agreement.

Step 5 Under Data Loss Prevention, select RSA Email DLP.

Step 6 Check the Enable RSA Email Data Loss Prevention check box.

Step 7 If message tracking is already enabled on your appliance, choose whether or not to enable matched content logging. By selecting this, the Cisco IronPort appliance logs DLP violations and AsyncOS displays the DLP violations and surrounding content in Message Tracking, including sensitive data such as credit card numbers and social security numbers.

Step 8 Submit and commit your changes.

Enabling RSA Enterprise Manager

If you want to use RSA Enterprise Manager to manage data loss prevention for your appliances, you need to configure your Email Security appliance as a partner device for Enterprise Manager. After you configure the RSA Enterprise Manager settings, the Email Security appliance sends its configuration to Enterprise Manager, which automatically adds the appliance as a partner device. The next time you open Enterprise Manager, the appliance will be shown as a partner device.

If you want to use SSL for communication between the Email Security appliance and Enterprise Manager, import one or more certificates to use as a server and client certificate into the appliance along with a certificate file for a certificate authority. The server and client certificates can be the same certificate, but must have the Email Security appliance's hostname for the common name. You can use a certificate generation tool provided by RSA to create the certificate, if you choose. See Certificates, page 11-28 for more information.

When you switch the Email Security appliance's data into RSA Enterprise Manager mode, the Email Security appliance saves your existing RSA Email DLP policies in case you switch back to RSA Email DLP mode later on.

Note See RSA's technical documentation on Enterprise Manager for information on managing DLP policies for the Email Security appliance.

Step 1 Select Security Services > RSA Email DLP.

Step 2 Click **Enable**.

Step 3 The license agreement page is displayed.

Note If you do not accept the license agreement, RSA Email DLP is not enabled on the appliance.

Step 4 Scroll to the bottom of the page and click **Accept** to accept the agreement.

Step 5 Under Data Loss Prevention, select RSA Enterprise Manager.

Step 6 Enter the hostname for the Enterprise Manager on your network that you want to use to manage DLP policies and `20000` for the port number. Separate the hostname and port number using a colon (`:`).

Step 7 Enter the service port on Email Security to which Enterprise Manager will connect.

Step 8 If you want the Email Security appliance and Enterprise Manager's connection to use SSL, check the Enable SSL Communication check box and then select the server certificate for Enterprise Manager and the client certificate for the Email Security appliance. The certificates must have the appliance's hostname for the common name. You can use the same certificate for both the client and server.

See Certificates, page 11-28 for information on setting up certificates for SSL communication between the appliance and Enterprise Manager.

Step 9 Choose whether to enable fingerprinting for source code and document detection If you select this option, Enterprise Manager sends fingerprinting detection content to the Email Security appliance. Fingerprinting can be used to detect the following:

- Databases
- Full or partial text matches in the text of a document
- Full binary match, which is a bit-by-bit exact match of a file

Step 10 If message tracking is already enabled on your appliance, choose whether or not to enable matched content logging. By selecting this, the Cisco IronPort appliance logs DLP violations and AsyncOS displays the DLP violations and surrounding content in Message Tracking, including sensitive data such as credit card numbers and social security numbers.

Step 11 Submit and commit your changes.

Exporting the DLP Configuration

If you want to use the active policies in your existing RSA Email DLP configuration to Enterprise Manager, you can export the configuration as a .zip file and import the policies into Enterprise Manager.

To create the .zip file, click **Export DLP Configuration** on the Data Loss Prevention Settings page. Enter a name for the .zip file and click **Export**. The Email Security appliance includes all active DLP policies assigned to an outgoing mail in the .zip file. Disabled DLP policies and DLP that are not assigned to an outgoing mail policy are not included in the .zip file.

If the Email Security appliance is part of the cluster, the appliance only exports the policies from the lowest level of the cluster. For example, if there are DLP policies at both the cluster and machine level, the appliance only exports the DLP policies from the machine level.

If the appliance is using RSA Enterprise Manager for DLP, you can use these instructions to export the active DLP policies that Enterprise Manager sent to the appliance.

The file is ready to be imported in Enterprise Manager. See the RSA Enterprise Manager help for instructions on importing the configuration into Enterprise Manager.

Switching Data Loss Prevention Modes

If you want to go back to using RSA Email DLP for data loss prevention after using RSA Enterprise Manager, use the Global Settings page to switch back to RSA Email DLP mode by following the steps in Enabling RSA Email DLP, page 11-3.

The Email Security appliance automatically reverts back to the RSA Email DLP policies it used before you configured it to use RSA Enterprise Manager mode. If the appliance did not use any local DLP policies when it was in RSA Email DLP mode, the appliance will continue to use the DLP policies from Enterprise Manager until you create a local DLP policy.

If you want to use local DLP policies similar to the ones on Enterprise Manager, you can recreate them using the DLP Policy Manager. The Email Security appliance does not automatically create new policies based on the ones used by Enterprise Manager and they cannot be imported from Enterprise Manager.

See Creating an Email DLP Policy Based on a Predefined Template, page 11-13 for information on creating DLP policies using the DLP Policy Manager.

See the RSA Enterprise Manager help for instructions on removing the Email Security appliance as a partner device in Enterprise Manager if you want to stop using Enterprise Manager to manage the appliance's DLP policies.

Message Actions

When the Email Security appliance detects a possible DLP violation in an outgoing message, it needs to know what to do with the message. Message actions define a *primary action* for the Email Security appliance to take with the message, which can be Deliver, Drop, or Quarantine. You can also specify *secondary actions* to take on messages. Secondary actions include:

- Sending a copy to a system quarantine if you choose to deliver the message. The copy is a perfect clone of the original, including the Message ID. Quarantining a copy allows you to test the RSA Email DLP system before deployment in addition to providing another way to monitor DLP violations. When you release the copy from the quarantine, the appliance delivers the copy to the recipient, who will have already received the original message.
- Encrypting messages. The appliance only encrypts the message body. It does not encrypt the message headers.
- Altering the subject header of messages containing a DLP violation.
- Adding disclaimer text to messages.

- Sending messages to an alternate destination mailhost.
- Sending copies (bcc) of messages to other recipients. (For example, you could copy messages with critical DLP violations to a compliance officer's mailbox for subsequent examination.)
- Sending a DLP violation notification message to the sender or other contacts, such as a manager or DLP compliance officer.

Message actions can be taken on all DLP policy severity levels except Ignore. See Setting the Severity Levels, page 11-15 for more information on severity levels for RSA Email DLP.

Note These actions are not mutually exclusive: you can combine some of them within different DLP policies for various processing needs for different user groups. You can also configure different treatments based on the different severity levels in the same policy. For example, you may want to quarantine messages with critical DLP violations and send a notification to a compliance officer but deliver messages with low severity levels.

For RSA Email DLP, specify the message actions you want your DLP policies to use when creating or editing the policies using the DLP Policy Manager. See DLP Policy Manager, page 11-11 for more information.

For RSA Enterprise Manager, create the message actions on your Email Security appliance first. The appliance sends the names and metadata of the message actions to Enterprise Manager, allowing you to use the actions in the DLP policies you create and manage in Enterprise Manager. See the RSA Enterprise Manager technical documentation for more information.

If you upgrade an appliance with existing DLP policies to AsyncOS 7.6, the operating system automatically converts the actions defined in the existing policies into message actions and updates the policies accordingly. AsyncOS generates names for the message actions but you can rename them using the DLP Message Actions page in the GUI. For information on renaming actions, see Editing a Message Action, page 11-8.

The DLP Message Actions page displays a list of the actions on your appliance. Click the **Policies** link in the Message Actions table to see the policies to which each action is assigned. Click the **Description** link to see a description of each action.

Figure 11-1 List of Actions on an Email Security Appliance

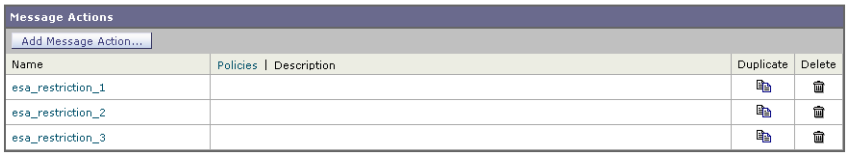

Creating a Message Action

Step 1 Select **Mail Policies > DLP Message Actions**.

Step 2 Click **Add Message Action**. The Add Message Action page is displayed.

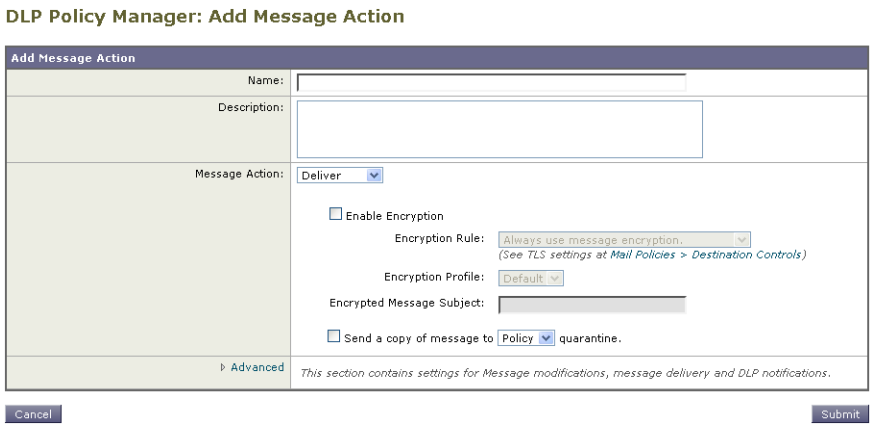

Step 3 Enter a name for the message action.

Step 4 Enter a description of the message action.

Step 5 Choose whether to drop, deliver, or quarantine messages containing DLP violations.

Note If you select Deliver, you can choose to have a copy of the message sent to a system quarantine. The copy of the message is a perfect clone, including the Message ID.

Step 6 If you want to encrypt the message upon delivery or its release from quarantine, check the Enable Encryption check box and select the following options:

- **Encryption Rule**. Always encrypts the message or only encrypt it if an attempt to send it over a TLS connection first fails.
- **Encryption Profile**. Encrypts the message using the specified encryption profile and delivers it if you use a Cisco IronPort Encryption Appliance or a hosted key service.
- **Encrypted Message Subject**. Subject for the encrypted message. Use the value is `$Subject` to keep the existing message subject.

Step 7 If you select Quarantine as the action, choose the quarantine you want to use for messages containing DLP violations.

Step 8 Click **Advanced** if you want to modify the message using any of the following options:

- Add a custom header
- Modify the message subject
- Deliver it to an alternate host
- Send a copy (bcc) to another recipient
- Send a DLP notification message

  For information on DLP notifications, see the "Text Resources" chapter in the *Cisco IronPort AsyncOS for Email Configuration Guide*.

Step 9 Submit and commit your changes.

Editing a Message Action

Step 1 Select **Security Settings > DLP Message Actions**.

Step 2 Click the name of the message action you want to edit.

Step 3 Modify the message action.

Step 4 Submit and commit your changes.

Deleting a Message Action

To delete a message action, click on the trash can icon next to the message action you want to delete. A confirmation message notifies you if the message is used in one or more DLP policies. Deleting a message action removes it from these DLP policies. Submit and commit your changes.

Duplicating a Message Action

If you want to create a message action that is similar to an existing one but with different settings, you have the option to create a duplicate message action.

Step 1 On the DLP Message Actions page, click on the **Duplicate** icon next to the message action that you want to duplicate.

Step 2 Enter a name for the new message action.

Step 3 Make your changes to the message action's settings.

Step 4 Submit and commit your changes.

RSA Email DLP

RSA Email DLP allows you to create and manage DLP policies locally on an individual Email Security appliance.

Understanding How RSA Email DLP Works

The RSA Email DLP feature uses a three-level policy structure to define your organization's data loss prevention rules for detecting sensitive data in outgoing messages:

- **Detection Rules.** At the lowest level, DLP content scanning consists of *detection rules* used to scan for particular patterns in a block of text. These detection rules include regular expressions, words and phrases, dictionaries, and entities, which are similar to the smart identifiers in used previously in AsyncOS.

- **Content Matching Classifier.** The next level is the *content matching classifier*, which scans an outgoing message and its attachments and headers for sensitive information, such as credit card data or other personal information. A classifier contains a number of detection rules along with *context rules* that impose additional requirements. As an example, consider the Credit Card Number classifier developed by RSA. This classifier not only requires that the message contains a text string that matches a credit card number pattern, but that it also contains supporting information such as an expiration date, a credit card company name (Visa, AMEX, etc.), or a person's name and address. Requiring this additional information results in more accurate verdicts of a message's content, leading to less false positives. A *DLP violation* occurs when a classifier detects sensitive information in a message.
- **DLP Policy.** At the highest level is a *DLP policy*, which consists of a set of conditions, as well as an assigned message action. The conditions include classifiers for a message's content and tests for message metadata, such as the sender, the recipient, or an attachment file type. The message action specifies both the overall action to take on messages (deliver, drop, or quarantine) and secondary actions such as encrypting the message, altering the header, and sending notifications to members of your organization.

You define your organization's DLP policies in the DLP Policy Manager and then enable the DLP policies in your outgoing mail policies. The appliance scans outgoing messages for DLP policy violations after the Outbreak Filters stage of the "work queue." AsyncOS also provides the DLP Assessment Wizard to guide you through setting up the most popular DLP policies. For more information, see Using the DLP Assessment Wizard, page 11-17.

The RSA Email DLP scanning engine scans each message, along with its headers and attachments, using every classifier in the DLP policies enabled in the outgoing mail policy. To scan message headers, the Cisco IronPort appliance's content scanning engine prepends the headers to the message body or any MIME parts that are content, and the RSA Email DLP scanning engine performs a content matching classifier scan. To scan attachments, the appliance's content scanning engine extracts the attachment for the RSA Email DLP scanning engine to analyze.

After scanning is complete, the RSA Email DLP engine determines if the message violated any of the enabled DLP policies. If the violation matches more than one DLP policy, the RSA Email DLP engine chooses the first matching DLP policy listed in the outgoing mail policy in a top-down fashion. You define the order of the DLP policies in the DLP Policy Manager.

The RSA Email DLP engine decides how to handle a message by first calculating a *risk factor* score for the DLP violation. The risk factor score represents the severity of the DLP violation, ranging from 0 to 100. The RSA Email DLP engine compares the risk factor score to the Severity Scale defined for that DLP policy. The Severity Scale categorizes the possible DLP violation as one of the following severity levels:

- Ignore
- Low
- Medium
- High
- Critical

The severity level determines which actions, if any, are taken on the message.

You can use the DLP Incidents report to view information on DLP violations discovered in outgoing mail. You can also use message tracking to search for messages based on the severity of the DLP violation.

- For more information on DLP email policies and content matching classifiers, see DLP Policies, page 11-10.

- For more information on content matching classifiers, see Content Matching Classifiers, page 11-20.
- For more information on the DLP Incidents report, see the "Using Email Security Monitor" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.
- For information on searching for messages with DLP violations in Message Tracking, see the "Tracking Email Messages" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.

Note The scanning engine only uses a classifier once when scanning a message. If an outgoing mail policy has two or more DLP policies that use the same classifier, the policies use the result from a single classifier scan.

Hardware Requirements

The RSA Email DLP feature is supported on all C-Series and X-Series appliances, except for the C10, C30, C60, C100, C300D, C350D, C360D, and C370D appliances.

DLP Policies

A DLP policy is a set of conditions that the RSA Email DLP scanning engine uses to determine whether an outgoing message contains sensitive data and the actions that AsyncOS takes when a message contains such data.

DLP policies include content matching classifiers developed by RSA, which the RSA Email DLP scanning engine uses to detect sensitive data in messages and attachments. The classifiers search for more than data patterns like credit card numbers and driver license IDs; they examine the context of the patterns leading to fewer false positives. For more information, see Content Matching Classifiers, page 11-20.

Before RSA Email DLP scanning takes place, AsyncOS's content scanning engine prepends the To, From, CC, and Subject headers to the message body, or any MIME parts that are tagged as content. This allows the RSA Email DLP scanning engine to scan these headers using the DLP policy's content matching classifiers.

If the DLP scanning engine detects a DLP violation in a message or an attachment, the DLP scanning engine determines the risk factor of the violation and returns the result to the matching DLP policy. The policy uses its own Severity Scale to evaluate the severity of the DLP violation based on the risk factor and applies the appropriate actions to the message. The scale includes five severity levels: Ignore, Low, Medium, High, and Critical. You decide what the Email Security appliance does with the message by specifying a message action for each severity level, except Ignore. For more information on message actions, see Message Actions, page 11-5.

Content of Policies

Email DLP policies contain the following information:

- Name and description of the policy.

- A list of content matching classifiers. Depending on the policy, you may be required to create a regular expression to search for identification numbers. See Content Matching Classifiers, page 11-20 for more information.
- A list of specific senders and recipients for filtering messages. See Filtering by Senders and Recipients, page 11-15 for more information.
- A list of attachment file types for filtering messages. See Filtering by Attachment Types, page 11-15 for more information.
- Severity settings, including the message actions applied to settings and Severity Scale adjustment. See Setting the Severity Levels, page 11-15 for more information.

DLP Policy Manager

The DLP Policy Manager is a single dashboard for managing all DLP policies on your Cisco IronPort appliance. You access the DLP Policy Manager from the Mail Policies menu. From the DLP Policy Manager you can perform the following actions:

- Create and manage DLP policies based on a predefined template. For more information, see Creating an Email DLP Policy Based on a Predefined Template, page 11-13.
- Create and manage DLP policies based on a custom template. For more information, see Creating a DLP Policy Using the Custom Policy Template, page 11-26.
- Create, import, and manage custom DLP dictionaries. For more information, see the "Text Resources" chapter in the *Cisco IronPort AsyncOS for Email Configuration Guide*.
- Manage US drivers license classifiers. For more information, see US Drivers License Classifiers, page 11-13.

Figure 11-2 DLP Policy Manager with Active DLP Policies

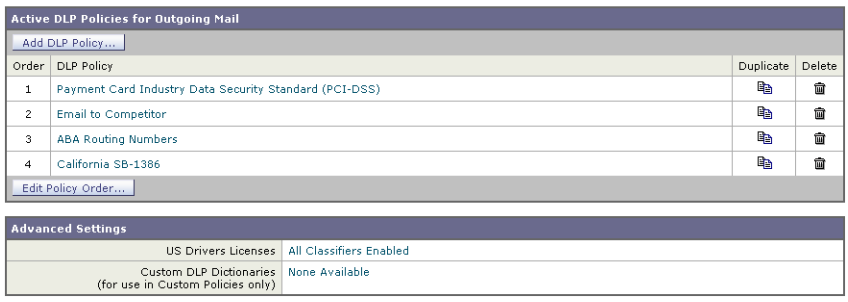

RSA Email DLP Policy Templates

AsyncOS comes with a large collection of predefined policy templates developed by RSA, Inc. to protect your organization's intellectual property and confidential information and enforce rules and regulations defined by laws and industry standards. When creating a DLP policy using the DLP Policy Manager, you first select the template that you want to use.

Figure 11-3 shows the categories of DLP policy templates available.

Figure 11-3 Add DLP Policy From Templates

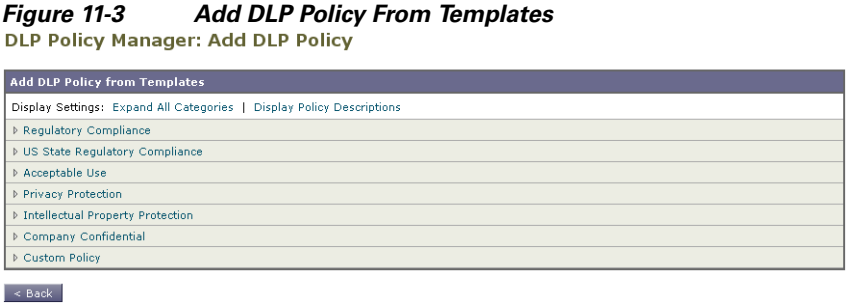

DLP policy templates are organized into the following categories:

- **Regulatory Compliance.** Identifies messages and attachments that contain personally identifiable information, credit information, or other protected or non-public information.
- **Acceptable Use.** Identifies messages sent to competitors or restricted recipients that contain sensitive information about an organization.
- **Privacy Protection.** Identifies messages and attachments that contain identification numbers for financial accounts, tax records, or national IDs.
- **Intellectual Property Protection.** Identifies popular publishing and design document file types that may contain intellectual property that an organization would want to protect.
- **Company Confidential.** Identifies documents and messages that contain information about corporate accounting information and upcoming mergers and acquisitions.
- **Custom Policy.** AsyncOS also provides the option to create your own policy from scratch using classifiers developed by RSA or your organization. This option is considered advanced and should be used only in the rare cases when the predefined policy templates do not meet the unique requirements of your network environment. See Advanced RSA Email DLP Policy Customization, page 11-25 for more information.

For information on DLP policy templates that require customization, see Customizing Classifiers for DLP Policies, page 11-14.

Figure 11-4 shows a predefined RSA policy template for detecting Payment Card Industry Data Security Standard (PCI-DSS) violations.

Figure 11-4 Predefined RSA Email DLP Policy Template

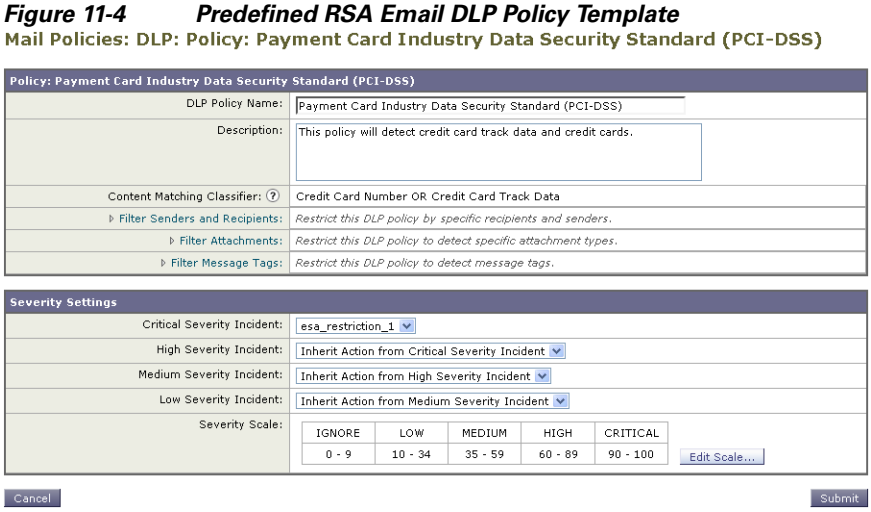

US Drivers License Classifiers

Many policies use a US Drivers License classifier. By default, this classifier searches for drivers licenses for all 50 US states and the District of Columbia. Even US state-specific policies such as California AB-1298 and Montana HB-732 search for all 51 types of US drivers licenses. If you are concerned about false positives or appliance performance, you can limit searching to specific US states or no states by clicking the link for US Drivers Licenses under **Advanced Settings** in the DLP Policy Manager. For more information on how the RSA scanning engine uses drivers license classifies, see US Drivers License, page 11-23.

Creating an Email DLP Policy Based on a Predefined Template

You can create a DLP policy using either a predefined template or a custom template. See Creating a DLP Policy Using the Custom Policy Template, page 11-26 for information on using a custom template.

Step 1 Select **Mail Policies > DLP Policy Manager**.

Step 2 Click **Add DLP Policy**.

Step 3 Click the name of a category to display a list of the available RSA Email DLP policy templates.

Note You can click **Display Policy Descriptions** to view detailed descriptions of the available policy templates.

Step 4 Click **Add** for the RSA Email DLP policy template that you want to use.

A page similar to Figure 11-4 on page 11-13 opens. All predefined templates already have a name and a description, which you can change. Most templates have one or more classifiers, and some have predefined attachment types for filtering messages.

Step 5 If the policy requires a customized classifier, enter a regular expression to define the pattern of your organization's identification numbering system and a list of words or phrases related to the identification numbers. See Customizing Classifiers for DLP Policies, page 11-14 for more information.

Note You cannot add or remove classifiers for policies based on a predefined template.

Step 6 Optionally, you can limit the DLP policy to messages with specific recipients or senders, attachment types, or message tags. For more information, see Filtering Messages for DLP Policies, page 11-14.

Step 7 In the Critical Severity Settings section, choose the action to perform on messages containing critical DLP violations.

Step 8 By default, the other severity levels inherit the message action from the level above it. If you want to define different settings for messages that match the high, medium, or low severity level, select the message action you want the appliance to perform.

Step 9 If you want adjust the DLP violation severity scale for the policy, click **Edit Scale** and adjust the settings. For more information, see Setting the Severity Levels, page 11-15.

Step 10 Submit and commit your changes.

The policy is added to the DLP Policy Manager.

Customizing Classifiers for DLP Policies

Some of the DLP policy templates require customized classifiers for better efficacy. These classifiers search for confidential identification numbers in outgoing messages, such as patient or student identification numbers, but require one or more regular expressions that define the patterns of your organization's record numbering system. You can also add a list of words and phrases that are associated with the record identification number for supporting information. If the classifier detects the number pattern in an outgoing message, it searches for the supporting information to verify that the pattern is an identification number and not a random number string. This results in less false positives.

For example, use the HIPAA and HITECH template to create a policy. This template includes the Patient Identification Numbers content matching classifier, which you can customize to detect a patient's identification number. Enter the regular expression `[0-9]{3}\-[A-Z]{2}[0-9]{6}` for the classifier. This regular expression detects numbers in the pattern of 123-CL456789. Enter "Patient ID" for a related phrase. Finish creating the policy and enable it in an outgoing mail policy. Submit and commit your changes. Now, if the policy detects the number pattern in an outgoing message with the phrase "Patient ID" in close proximity to the number pattern, the DLP policy returns a DLP violation.

For information on how to create a regular expression, see Regular Expressions for Content Matching Classifiers, page 11-24. For more information on how content matching classifiers detect DLP violations, see Content Matching Classifiers, page 11-20.

Filtering Messages for DLP Policies

You have the option of limiting a DLP policy to scanning only messages based on specific information first detected by AsyncOS. DLP policy scanning can be limited by the following information:

- Senders and recipients
- Attachment types
- Message tags

Filtering by Senders and Recipients

You can limit the DLP policy to scan messages with specific recipients or senders in one of the following ways:

- Full email address: `user@example.com`
- Partial email address: `user@`
- All users in a domain: `@example.com`
- All users in a partial domain: `@.example.com`

You can separate multiple entries using a line break or a comma.

For an outgoing message, AsyncOS first matches the recipient or sender to an outgoing mail policy. After the recipient or sender is matched, RSA Email DLP then matches the sender or recipient to the DLP policies enabled for the mail policy.

Filtering by Attachment Types

You can limit the DLP policy to messages with specific attachment types. Attachments are first extracted using AsyncOS's content scanning engine and then the content of the attachment is scanned by the RSA Email DLP scanning engine. The appliance provides a number of predefined file types for scanning, but you can also specify file types that are not listed. If you specify a file type that is not predefined, AsyncOS searches for the file type based on the attachment's extension. You can limit RSA Email DLP scanning to attachments with a minimum file size in bytes.

Filtering by Message Tag

If you want to limit a DLP policy to scanning messages containing a specific phrase, you can use a message or content filter to search outgoing messages for the phrase and insert a custom message tag into the message. When creating a DLP policy, select the message tags you want to use for filtering outgoing messages. For more information, see Content Filter Actions, page 6-12 and the "Using Message Filters to Enforce Mail Policies" in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Setting the Severity Levels

If RSA Email DLP scanning engine detects a DLP violation, it calculates a risk factor score that represents the severity of the violation, ranging from 0 to 100. The policy compares the risk factor score to the Severity Scale. The Severity Scale includes five severity levels: Ignore, Low, Medium, High, and Critical. The severity level determines the actions taken on the message. By default, all severity levels (except Ignore) inherit the settings of the higher severity level; the High severity level inherits the settings from Critical, Medium inherits from High, and Low inherits from Medium. You can edit the level to specify different actions for different severities.

For information on how the DLP scanning engine calculates a risk factor, see Understanding How RSA Email DLP Works, page 11-8.

You can also adjust the Severity Scale for a policy to define the estimated severity of the DLP violation returned by the scanning engine. Figure 11-5 shows the severity scale. Use the scale's arrows to adjust the scores for the severity levels.

Figure 11-5 Adjusting the DLP Policy Severity Scale

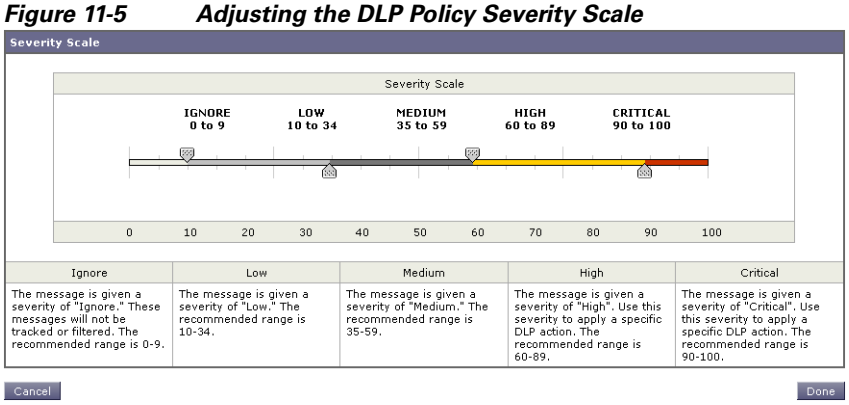

Arranging the Order of the Email DLP Policies

The order of policies in the DLP Policy Manager is important. When a DLP violation occurs, RSA Email DLP matches the violation to the DLP policies enabled in the outgoing mail policy. If the violation matches more than one DLP policy, RSA Email DLP chooses the first matching DLP policy in the top-down order.

Step 1 On the DLP Policy Manager page, click **Edit Policy Order**.

Step 2 Click on the row for a policy you want to move and drag it to a new position in the order.

Step 3 Once you have finished reordering the policies, submit and commit your changes.

Editing an Email DLP Policy

Step 1 Click on the name of the RSA Email DLP policy in the listing on the DLP Policy Manager page. The Mail Policies: DLP page is displayed.

Step 2 Make changes to the DLP policy.

Step 3 Submit and commit your changes.

Note If you rename a policy, you will have to re-enable it in your outgoing mail policies.

Deleting an Email DLP Policy

To delete a DLP policy, click on the trash can icon next to the policy in the listing. A confirmation message is displayed. This message notifies you if the DLP policy is used in one or more outgoing mail policies. Deleting a DLP policy removes it from these mail policies. Submit and commit your changes.

Duplicating an Email DLP Policy

If you want to create a DLP policy that is similar to an existing one but with some different settings, the DLP Policy Manager gives you the option to create a duplicate policy.

Step 1 On the DLP Policy Manager page, click on the **Duplicate** icon next to the policy in the listing that you want to duplicate.

Step 2 Enter a name for the policy.

Step 3 Make your changes to the policy's settings.

Step 4 Submit and commit your changes.

Using the DLP Assessment Wizard

AsyncOS provides a browser-based DLP Assessment Wizard to guide you through the three-step process of configuring popular RSA Email DLP policies and enabling them in the appliance's default outgoing mail policy.

The DLP policies added using the DLP Assessment Wizard deliver all messages by default, regardless of the severity of detected DLP violations. Use the DLP Policy Manager to edit the overall action on messages, the severity level settings, and any other policy settings you want. For more information on editing DLP policies, see DLP Policy Manager, page 11-11.

To launch the DLP Assessment Wizard when the appliance is using RSA Email DLP, go to the **Security Services > RSA Email DLP** page. Click **Edit Settings**. Check the **Enable and configure DLP using the DLP Assessment Wizard** check box and click **Enable**.

You can only use the DLP Assessment Wizard if there are no existing DLP policies on the appliance.

Figure 11-6 shows the option of running the DLP Assessment Wizard from the RSA Email Data Loss Prevention Settings page.

Figure 11-6 RSA Email Data Loss Prevention Settings Page

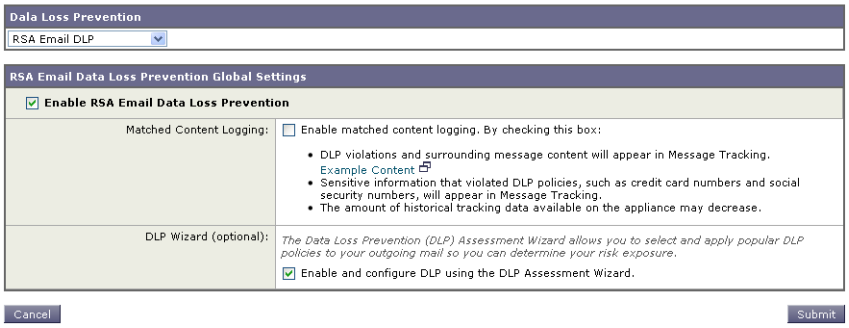

Running the DLP Assessment Wizard

The DLP Assessment Wizard walks you through completing the following DLP configuration tasks, broken down into three steps:

Step 1 Policies

- Select the DLP policies for the types of information you want to protect on your network
- Customize the DLP policies that require additional information to find sensitive data

Step 2 Reports

- Configure DLP Incident Summary report delivery settings

Step 3 Review

- Review and enable your DLP policies

Step through the DLP Assessment Wizard, clicking **Next** after you complete each step. You can move back to a previous step by clicking **Previous**. At the end of the process, you are prompted to commit the changes you have made. Your changes will not take effect until they have been committed.

Step 1: Policies

Selecting the DLP Policies

Select the DLP policies for the types of sensitive information you want the appliance to detect in outgoing messages.

The following policies are available:

- **FERPA (Family Educational Rights and Privacy Act)** detects student records and can be customized to detect student identification numbers.
- **GLBA (Gramm-Leach Bliley Act)** detects credit card numbers, US Social Security numbers, US Drivers License numbers and may be customized to detect custom account numbers.

- **California SB-1386** detects documents and transmissions that contain personally identifiable information (PII) as regulated by California SB-1386 (Civil Code 1798), such as US Social Security numbers, credit card numbers, and US drivers license numbers. Any business that operates in California and owns or licenses computerized PII data for California residents, regardless of their physical location, is required to comply.
- **Restricted Files** detects emails that contain restricted files, including .mdb, .exe, .bat and Oracle executable files (.fmx, .frm). This policy can be customized to add additional file attributes to the policy violation rules.

You can create other types of DLP policies using the DLP Policy Manager.

Customizing the DLP Policies

Some DLP policies use content matching classifiers that can be customized to detect sensitive information in outgoing messages. The customized classifiers for the FERPA and GLB, policies use a regular expression to search for identification number patterns in outgoing messages. If you select the Restricted Files policy, you can choose the attachment file types you want the DLP policy to detect. The Restricted Files policy detects .exe and .mdb files by default, but you can remove these file types. You can also configure the Restricted Files policy to apply only to encrypted or password-protected files.

For more information on customizing the content matching classifiers for these DLP policies, see Customizing Classifiers for DLP Policies, page 11-14.

Click **Next** to continue.

Figure 11-7 DLP Assessment Wizard: Step 1. Policies

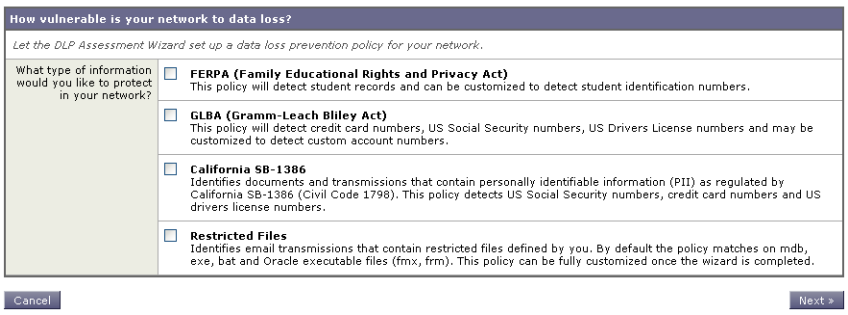

Step 2: Reports

Enter an email address for the scheduled DLP Incident Summary report. Use commas to separate multiple addresses. If you leave this value blank, the scheduled report is not created. For more information on DLP Incident Summary reports, see the “Using Email Security Monitor” chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.

Click **Next** to continue.

Figure 11-8 DLP Assessment Wizard: Step 2. Reports

Step 3: Review

A summary of the DLP configuration information is displayed. You can edit the Policies and Reporting 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.

Figure 11-9 DLP Assessment Wizard: Step 3. Review

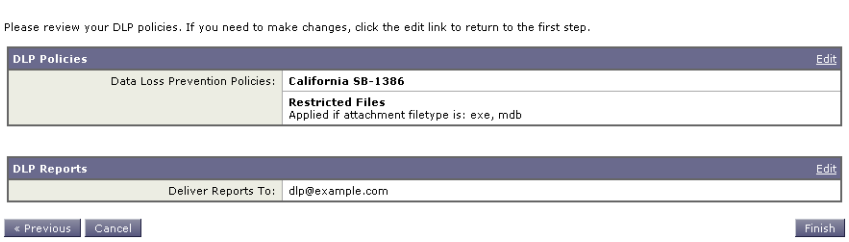

Once you are satisfied with the information displayed click **Finish**. AsyncOS displays the Outgoing Mail Policies page with your DLP policies enabled in the default outgoing mail policy. A summary of your DLP policy configuration is displayed at the top of the page. Commit your changes.

For information on editing the DLP policies and creating additional ones, see DLP Policy Manager, page 11-11. For information on enabling the DLP policies for other outgoing mail policies, see Configuring Per-Recipient Policies for DLP, page 11-31.

Content Matching Classifiers

Content matching classifiers are the detection components of the RSA Email DLP scanning engine. They search messages, message headers, and the content of extracted attachments for data patterns, such as credit card numbers or driver license identification numbers, and the context in which the patterns appear. For example, a classifier for detecting credit card numbers scans for not only patterns of numbers that match the credit card number format, but supporting data like expiration dates and the names of credit card companies. Evaluating the context of the data decreases the number of false positives.

Many of the policy templates from RSA include a predefined set of classifiers. When creating a policy based on the Custom Policy template, you can choose an RSA classifier or add one of your own. For information on creating your own classifier to use in custom DLP policies, see Creating a Content Matching Classifier, page 11-26.

A number of policy templates require customization of one or more classifiers in order to detect sensitive data. Customization includes creating a regular expression to search for identification numbers and a list of words and phrases that may consistently appear with the identification number. For example, adding a policy based on the FERPA (Family Educational Rights and Privacy Act) template requires creating a regular expression to match custom student ID numbers. If the ID numbers consistently appear with the phrase “Student ID,” such as “Student ID: 123-45-6789,” adding the phrase to the policy would improve content matching accuracy. For more information on required customization for DLP policies, see Customizing Classifiers for DLP Policies, page 11-14.

Note For policies that do not have a classifier, the scanning engine always returns a risk factor value of "75" when a message violates the policy. You may want to adjust the severity scale for such policies, depending on the type of DLP violations that may occur. See Setting the Severity Levels, page 11-15 for more information.

Classifier Detection Rules

Classifiers require rules for detecting DLP violations in a message or document. Classifiers can use one or more of the following detection rules:

- **Words or Phrases**. A list of words and phrases for which the classifier should look. Separate multiple entries with a comma or line break.
- **Regular Expression**. A regular expression to define a search pattern for a message or attachment. You can also define a pattern to exclude from matching to prevent false positives. See Examples of Regular Expressions for DLP, page 11-25 for more information.
- **Dictionary**. A dictionary of related words and phrases. RSA Email DLP comes with dictionaries created by RSA, but you can create your own. See the Chapter 14, "Text Resources" for more information.
- **Entity**. Similar to smart identifiers in previous versions of AsyncOS, an entity identifies patterns in data, such as ABA routing numbers, credit card numbers, addresses, and social security numbers.

Classifiers assign a numeric value to the detection rule matches found in a message and calculate a score for the message. The risk factor used to determine the severity of a message's DLP violation is a 0 - 100 version of the classifier's final score. Classifiers use the following values to detect patterns and calculate the risk factor:

- **Proximity**. How close the rule matches must occur in the message or attachment to count as valid. For example, if a numeric pattern similar to a social security number appears near the top of a long message and an address appears in the sender's signature at the bottom, they are probably not related and the classifier does not count them as a match.
- **Minimum Total Score**. The minimum score required for the classifier to return a result. If the score of a message's matches does not meet the minimum total score, its data is not considered sensitive.
- **Weight**. For each rule, you specify a "weight" to indicate the importance of the rule. The classifier scores the message by multiplying the number of detection rule matches by the weight of the rule. Two instances of a rule with a weight of `10` results in a score of `20`. If one rule is more important for the classifier than the others, it should be assigned a greater weight.
- **Maximum Score**. A rule's maximum score prevents a large number of matches for a low-weight rule to skew the final score of the scan.

To calculate the risk factor, the classifier multiplies the number of matches for a detection rule by the weight of the rule. If this value exceeds the detection rule's maximum score, the classifier uses the maximum score value. If the classifier has more than one detection rule, it adds the scores for all of its detection rules into a single value. The classifier maps the detection rules score (10 - 10000) on a scale of 10 -100 using the logarithmic scale shown in Table 11-1 to create the risk factor.

Table 11-1 Logarithmic Scale for Calculating the Risk Factor

Rule Scores	Risk Factor
10	10
20	20

Table 11-1 Logarithmic Scale for Calculating the Risk Factor

Rule Scores	Risk Factor
30	30
50	40
100	50
150	60
300	70
500	80
1000	90
10000	100

Classifier Examples

The following examples show how classifiers match message content.

Credit Card Number

Several DLP policy templates include the Credit Card Number classifier. The credit card number itself is subject to various constraints, such as the pattern of digits and punctuation, the issuer-specific prefix, and the final check digit. The classifier requires additional supporting information to make a match, such as a second credit card number, an expiration date, or the name of the card issuer. This reduces the number of false positives.

Examples:

- `4999-9999-9999-9996` (No match because of no supporting information)
- `4999-9999-9999-9996 01/09` (Match)
- `Visa 4999-9999-9999-9996` (Match)
- `4999-9999-9999-9996 4899 9999 9999 9997` (Match because of more than one credit card number)

US Social Security Number

The US Social Security Number classifier requires a properly formatted number as well as supporting data, such as a date of birth, name, or the string `SSN`.

Examples:

- `321-02-3456` (No match because of no supporting information)
- `321-02-3456 July 4` (Match)
- `321-02-3456 7/4/1980` (Match)
- `321-02-3456 7/4` (No match)
- `321-02-3456 321-02-7654` (Match because of more than one SSN)
- `SSN: 321-02-3456` (Match)
- `Joe Smith 321-02-3456` (Match)
- `321-02-3456 CA 94066` (Match)

ABA Routing Number

The ABA Routing Number classifier is similar to the Credit Card Number classifier.

Examples:

- `119999992` (No match because of no supporting information)
- `routing 119999992 account 1234567` (Match)

US Drivers License

Several DLP policy templates use the US Drivers License classifier. This classifier contains a separate set of detection rules for each US state and the District of Columbia. You can selectively enable or disable states that are not important for your organization's policies by clicking the link for US Drivers Licenses under **Advanced Settings** in the DLP Policy Manager.

Note A predefined DLP policy template for a specific state, such as California SB 1386, uses the detection rules for all states and will return a DLP violation for data with a non-California driver license because it is still considered a privacy violation.

The individual state classifiers match against the patterns for that state, and requires the corresponding state name or abbreviation, and additional supporting data.

Examples:

- `CA DL: C3452362` (Match because it has the correct pattern for the number and supporting data)
- `California DL: C3452362` (Match)
- `DL: C3452362` (No match because there is not enough supporting data)
- `California C3452362` (No match because there is not enough supporting data)
- `OR DL: C3452362` (No match because it is the incorrect pattern for Oregon)
- `OR DL: 3452362` (Match because it is the correct pattern for Oregon)
- `WV DL: D654321` (Match because it is the correct pattern for West Virginia)
- `WV DL: G6543` (No match because it is the incorrect pattern for West Virginia)

US National Provider Identifier

The US National Provider Identifier classifier scans for a US National Provider Identifier (NPI) numbers, which is a 10-digit number with a check digit.

Examples:

- `NPI: 3459872347` (Match for NPI)
- `3459872347` (No match because of no supporting information)
- `NPI: 3459872342` (No match because of incorrect check digit)

Student Records

The predefined FERPA (Family Educational Rights and Privacy Act) DLP policy template uses the Student Records classifier. Combine it with a customized Student Identification Number classifier to detect specific student ID patterns for better accuracy.

Example:

- `Joe Smith, Class Rank: 234, Major: Chemistry Transcript` (Match)

Corporate Financials

The predefined Sarbanes-Oxley (SOX) policy template uses the Corporate Financials classifier to search for non-public corporate financial information.

Examples:

`2009 Cisco net sales, net income, depreciation` (Match)

`FORM 10-Q 2009 I.R.S. Employer Identification No.` (Match)

Regular Expressions for Content Matching Classifiers

A number of policy templates require customization of one or more classifiers, which involves creating a regular expression to search for identification numbers that may be linked to confidential information, such as a custom account number or patient identification number. The style of regular expressions used for content matching classifiers is the **POSIX Basic Regular Expression** style regular expressions.

Use the following table as a guide for creating regular expressions for classifiers:

Table 11-2 Regular Expression in Classifiers

Regular expression (`abc`)	Regular expressions for classifiers match a string if the sequence of directives in the regular expression match any part of the string. For example, the regular expression `ACC` matches the string `ACCOUNT` as well as `ACCT`.
[ ]	Use brackets to indicate a set of characters. Characters can defined individually or within a range. For example, `[a-z]` matches all lowercase letters from `a` to `z`, while `[a-zA-Z]` matches all uppercase and lowercase letters from `A` to `Z`. `[xyz]` matches only the letters `x`, `y`, or `z`.
Backslash special characters (\)	The backslash character *escapes* special characters. Thus the sequence `\.` only matches a literal period, the sequence `\$` only matches a literal dollar sign, and the sequence `\^` only matches a literal caret symbol. The backslash character also begins tokens, such as `\d`. **Important Note:** The backslash is also a special escape character for the parser. As a result, if you want to include a backslash in your regular expression, you must use *two* backslashes — so that after parsing, only one "real" backslash remains, which is then passed to the regular expression system.
\d	Token that matches a digit (`0-9`). To match more than one digit, enter an integer in `{}` to define the length of the number. For example, `\d` matches only a single digit such as `5`, but not `55`. Using `\d{2}` matches a number consisting of two digits, such as `55`, but not `5`.

Table 11-2 Regular Expression in Classifiers

Number of repetitions `{min,max}`	The regular expression notation that indicates the number of repetitions of the previous token is supported. For example, the expression "`\d{8}`" matches `12345678` and `11223344` but not `8`.
Or (|)	Alternation, or the "or" operator. If A and B are regular expressions, the expression "`A|B`" will match any string that matches either "A" or "B." Can be used to combine number patterns in a regular expression. For example, the expression "`foo|bar`" will match either `foo` or `bar`, but not `foobar`.

Examples of Regular Expressions for DLP

The primary case for using regular expressions in content matching classifiers is to detect specific account, patient, or student identification numbers. These are usually simple regular expressions that describe patterns of numbers and letters. For example:

- An 8-digit number: `\d{8}`
- Identification code with hyphens between sets of numbers: `\d{3}-\d{4}-\d`
- Identification code that begins with a single letter that can be upper or lower case: `[a-zA-Z]\d{7}`
- Identification code that begins with three digits and is followed by nine uppercase letters: `\d{3}[A-Z]{9}`
- Using `|` to define two different number patterns to search for: `\d{3}[A-Z]{9}|\d{2}[A-Z]{9}-\d`

Note Regular expressions are case sensitive, so they should include upper and lower case, such as `[a-zA-Z]`. If only certain letters are used, you can define the regular expression accordingly.

The less specific the pattern, such as an 8-digit number, the more likely you will want the policy to search for additional words and phrases to distinguish a random 8-digit number from an actual customer number.

Advanced RSA Email DLP Policy Customization

If the available RSA Email DLP policy templates do not meet the unique requirements of your organization, a number of options are available for creating your own DLP policies from scratch. These options include:

- Creating your own DLP policy using the Custom Policy Template
- Creating your own classifiers to use in a custom policy
- Creating and importing your own DLP dictionaries to use in a custom policy

Note These options are advanced and should only be used in cases where predefined settings do not meet your organization's needs.

Creating a DLP Policy Using the Custom Policy Template

You can create a custom DLP policy using the Custom Policy template add either a predefined RSA classifier or a custom classifier to the policy. See Creating a Content Matching Classifier, page 11-26 for instructions on creating a classifier.

Custom policies can return a DLP violation if the content matches a single classifier or all classifiers, depending on how the policy is defined. To prevent false positives, a DLP policy can include a classifier that the message content must not match. By checking the NOT checkbox for a classifier, a message that includes matching content for the classifier is not reported as a DLP violation.

Step 1 Select **Mail Policies > DLP Policy Manager**.

Step 2 Click **Add DLP Policy**.

Step 3 Click the name of the Custom Policy category.

Step 4 Click **Add** for the Custom Policy template.

Step 5 Enter a name and description for the policy.

Step 6 Select a classifier for the policy. You can use an existing classifier or select the option **Create a Classifier**.

Step 7 Click **Add**.

If you selected **Create a Classifier**, the Add Content Matching Classifier page opens. Otherwise, the predefined classifier is added to the policy.

Step 8 To add more than one classifier to the policy, repeat steps 6 - 7.

Step 9 Optionally, you can limit the DLP policy to messages with specific recipients, senders, or attachment types. You can separate multiple entries using a line break or a comma. For more information, see Filtering by Senders and Recipients, page 11-15 and Filtering by Attachment Types, page 11-15.

Step 10 In the Critical Violations Settings section, choose the action to take on messages containing critical DLP violations.

Step 11 If you want to define different settings for messages that match the high, medium, or low severity level, uncheck the **Inherit settings** check box for the appropriate security level. Edit the overall action for the message and the other settings.

Step 12 If you want to adjust the DLP violation severity scale for the policy, click **Edit Scale** and adjust the settings. For more information, see Setting the Severity Levels, page 11-15.

Step 13 Submit and commit your changes.

The policy is added to the DLP Policy Manager.

Creating a Content Matching Classifier

When creating a custom policy, you can create a custom classifier by selecting the **Create a Classifier** option. See Classifier Detection Rules, page 11-21 for more information on the rules and values required to create a classifier.

After you have created and submitted the classifier, it will appear in the list of available classifiers when creating a custom policy.

Step 1 Enter a name and description for the classifier.

Step 2 Enter the number of characters within which the classifier's rules must be found in order to count as a proximity match.

Step 3 Enter the minimum total score for the classifier.

Step 4 Define a rule for the classifier, including the weight and maximum score.

Step 5 Click **Add Rule** to add the rule to the classifier. You can add multiple rules.

Step 6 Submit your classifier and continue creating the custom policy.

RSA Enterprise Manager

Starting in AsyncOS 7.6, Cisco provides the option of using RSA Enterprise Manager to create and manage DLP policies for the Email Security appliances on your network. RSA Enterprise Manager is a third-party software offered by RSA Security, Inc. It is not a part of the Cisco IronPort Email Security appliance. Partnering your Email Security appliances with Enterprise Manager opens up a robust set of DLP capabilities to your appliances while turning management of the appliances' DLP functionality over to Enterprise Manager. RSA Enterprise Manager also acts as a centralized manager for DLP across all connected Email Security appliances.

Note This guide provides an overview of how the Email Security appliance integrates with RSA Enterprise Manager along with instructions on configuring the Email Security appliance. Use the RSA Enterprise Manager technical documentation for information on configuring Enterprise Manager to work with the Email Security appliance and managing DLP policies using Enterprise Manager. This guide references the RSA Enterprise Manager help when appropriate.

How RSA Enterprise Manager DLP Works

When you use RSA Enterprise Manager for DLP, Enterprise Manager becomes your interface for managing the DLP policies for the Email Security appliances on your network and handling messages that contain DLP violations. Setting up RSA Enterprise Manager for DLP requires you to configure both your Email Security appliances and Enterprise Manager to work together to exchange data.

First, create the outgoing mail policies and message actions you want to use for your DLP monitoring and enforcement on the Email Security appliance. When you connect the Email Security appliance to Enterprise Manager, as described in Enabling RSA Enterprise Manager, page 11-3, the Email Security appliance sends the names and metadata of the mail policies and message actions to Enterprise Manager, where you use this information when creating DLP policies. After you create and enable a DLP policy in Enterprise Manager, Enterprise Manager sends the DLP policy as part of a data package to the Email Security appliance. The Email Security appliance stores the DLP policies and uses them to scan outgoing messages for violations based on the Enterprise Manager DLP policies and sends information on incidents of DLP violations to Enterprise Manager for the administrator to view and manage. RSA Enterprise Manager requires the User Distinguished Name LDAP query to retrieve the sender's name from messages in order to include this information as part of the DLP incident data sent by the appliance when it detects a DLP violation.

The order of the DLP policies defined in Enterprise Manager is important. When a DLP violation occurs, the Email Security appliance matches DLP violations in a top-down manner and takes action against the message based on the first policy it matches. You configure the policy order in Enterprise Manager, which is sent as part of the data package to the Email Security appliance.

Use the **Security Services > RSA Email DLP** page to see information on the latest DLP policy updates from Enterprise Manager and the Mail Policies > DLP Policy Manager page to enable and disable individual DLP policies for the Email Security appliance.

The Email Security appliance will continue to use any existing local RSA Email DLP policies until it receives its first package of DLP policies from Enterprise Manager.

Setting Up the Email Security Appliance for RSA Enterprise Manager DLP

There are a number of settings on the Email Security appliance that you need to configure in order for Enterprise Manager to work with the Email Security appliance.

Certificates

If you want to use an SSL connection between the Email Security appliance and Enterprise Manager, you will need a one or more certificates and signing keys from a recognized certificate authority to use for mutual authorization of the two machines. The common name of the certificates should be the appliance's hostname. Use the Email Security appliance's **Networks > Certificates** page to manage the certificates and add the certificate authority to the appliance's list of recognized certificate authorities. When configuring an SSL connection using the DLP Global Settings, the Enterprise Manager server is the server and the Email Security appliance is the client.

RSA Enterprise Manager provides a certificate generation tool that you can use to generate a .p12 file that you can use as both the server and client certificate for the connection. This tool can only generate a single certificate. If you want to use different certificates for the appliance and the Enterprise Manager server, you will have to get them from another source.

The directory on the Enterprise Manager server that contains the .p12 certificate file also has a .pem certificate file. Import this file onto the Email Security appliance as a certificate authority if you want to use the .p12 file.

Step 1 Open a command prompt.

Step 2 Change to `C:\Program Files\RSA\Enterprise Manager\etc`.

Step 3 Run the following command:

```
"%JAVA_HOME%/bin/java" -cp ./emcerttool.jar
com.rsa.dlp.tem.X509CertGenerator -clientservercasigned -cacn <NAME OF CAPROVIDED DURING INSTALL> -cakeystore catem-keystore -castorepass <PASSWORD FOR CA PROVIDED DURING INSTALL>  -cn <DEVICE_CN> -storepass <DEVICE STORE PASSWORD> -keystore <NAME OF DEVICE STORE>
```

A sample command may look like the following:

```
"%JAVA_HOME%/bin/java" -cp ./emcerttool.jar
com.rsa.dlp.tem.X509CertGenerator -clientservercasigned -cacn emc-cisco
-cakeystore catem-keystore -castorepass esaem -cn ironport -storepass esaem
-keystore device-store
```

This outputs the `device-store.p12` file in the same folder.

Step 4 `<NAME OF DEVICE STORE>.p12` is the desired file. Please use this on Email Security appliance as its certificate.

You can also use the following additional command-line switches:

```
-org <value in double quotes if it contains space>

-orgunit <value in double quotes if it contains space>

-title <value in double quotes if it contains space>

-validity <number of days>
```

To use client and server certificates for the SSL connection:

1. Add the certificate authority to the appliance using the **Networks > Certificates** page. If you generated a client/server certificate using the tool provided by RSA, import the .pem file.
1. Upload the client and server certificate(s) to the Email Security appliance using the **Networks > Certificates** page. See the "Customizing Listeners" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information. You can use the same certificate for the client and server. If you generated a certificate using the RSA tool, import the .p12 certificate and use it for both the client and server certificate.
2. The common name of the client and server certificates must be the hostname of the Email Security appliance.
3. When configuring the SSL connection using the DLP Global Settings, assign the client certificate to the Email Security appliance and the server certificate to Enterprise Manager. See Data Loss Prevention Global Settings, page 11-2 for more information.

If Enterprise Manager manages the connected Email Security appliances at the group or cluster level, the appliances should each have their own certificate with a common name that matches their appliance's hostname, but all of the certificates should have the same certificate name. Use the Network > Certificates page on the appliances to make sure that the certificate names match. If a certificate cannot be found on the Email Security appliance, Enterprise Manager disconnects the appliance.

LDAP User Distinguished Name Query

When the Email Security appliance sends data to Enterprise Manager on DLP incident, the appliance must include the complete distinguished names for the message senders. To acquire the sender name for Enterprise Manager, create a user distinguished name query for your LDAP server and add the query to the listeners that send outgoing messages on your Email Security appliance. The Email Security appliance only uses this query when RSA Enterprise Manager is enabled for DLP. See the "LDAP Queries" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information.

Message Actions

When you create message actions on the Email Security appliance, the appliance sends the name of the action and some read-only metadata about the action to Enterprise Manager for DLP policies. You cannot use Enterprise Manager to modify the action or create new ones.

Message actions can order the Email Security appliance to notify a user, such as a DLP compliance officer, if a DLP violation occurs. Enterprise Manager's DLP policies can also send DLP violation notifications to users. Cisco recommends that you set up notifications using either Enterprise Manager or the Message Actions page in the Email Security appliance, but not both, to prevent duplicate notifications.

See Chapter 11, "Message Actions" for more information.

DLP Policy Manager for Enterprise Manager DLP Policies

The DLP Policy Manager shows the RSA Enterprise Manager DLP policies currently in use on the Email Security appliance. You can use the Manager to enable or disable individual DLP policies on the Email Security appliance. Any outgoing mail policies assigned to the disable DLP policy will skip the policy when evaluating messages for DLP violations.

Figure 11-10 Enterprise Manager DLP Policies in DLP Policy Manager

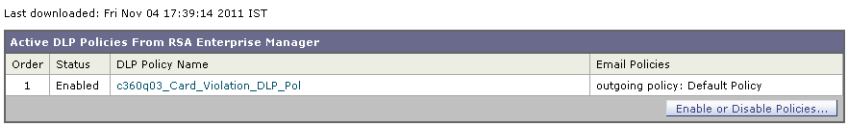

If the Email Security appliance has not received the DLP policies from Enterprise Manager, it will continue to use any existing RSA Email DLP policies until it receives a data package with the new policies from Enterprise Manager.

RSA Enterprise Manager and Language Support

The Email Security appliance displays any data it receives from RSA Enterprise Manager in the language that was used in Enterprise Manager. The appliance does not display this information in the language you selected for the appliance interface. This applies to DLP policies, classifiers, dictionaries, and anything else created in Enterprise Manager that the appliance receives in the data package. For example, if the DLP policies and classifiers from Enterprise Manager were written in English but the interface of the Email Security appliance is displayed in French, the Email Security appliance displays the name and descriptions of the DLP policies and classifiers from Enterprise Manager in English. The rest of the interface remains in French.

Quarantines

If a message containing a DLP violation matches a DLP policy that requires the message to be quarantined, the Email Security appliance sends the message to the quarantine specified by the DLP policy's message action. The user responsible for evaluating DLP violations can review the incident using Enterprise Manager and can then use Enterprise Manager to instruct the appliance to release or delete the message from the quarantine. If the message action requires the message to be encrypted on release, it is the Email Security appliance that encrypts the message, not Enterprise Manager.

Users can view messages quarantined by Enterprise Manager using the **Monitor > Quarantines** page in the Email Security appliance's GUI. Cisco recommends that users only release or delete messages with DLP violations from Enterprise Manager, not the local Email Security appliance's GUI.

Cisco also recommends the following procedures for using quarantines with Enterprise Manager:

- Use one or more dedicated quarantines for DLP violations.
- Set a timeout large enough for Enterprise Manager to complete its tasks.
- Be aware that Email Security appliance will still release or delete quarantine messages when the quarantine exceeds the allotted space.

For more information on how quarantines work on the Email Security appliance, see the "Quarantines" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.

Connectivity Between the Email Security Appliance and Enterprise Manager

In cases where connectivity between the Email Security appliance and Enterprise Manger is lost, any data that the appliance and Enterprise Manager cannot send is queued for delivery until the connection is restored. For the appliance, that means any data on messages containing possible DLP violations is queued. For Enterprise Manager, that means any data packages with new DLP policy information are queued. In cases where the appliance does not receive updated DLP policy data from Enterprise Manager, the appliance continues to use the DLP policies it had previously received from Enterprise Manager.

Using Enterprise Manager with Clustered Appliances

If you are using Enterprise Manager to manage the DLP policies for clustered Email Security appliances, be aware of the following:

- The Email Security appliance sends Enterprise Manager the outgoing mail policies and message actions from the lowest cluster level where these settings are configured. If these settings are configured differently at the cluster and machine level, the Email Security appliance sends Enterprise Manager the settings from the machine level. If you want to use the outgoing mail policies and message actions configured at a higher cluster level, delete the policies and actions defined at the lower levels that you do not want to use.
- The Email Security appliance uses the Data Loss Prevention mode used at the lowest cluster level where this setting is configured. For example, if a clustered appliance is configured to use the local RSA Email DLP mode at machine level and RSA Enterprise Manager at the cluster level, the appliance uses RSA Email DLP for data loss prevention and does not communicate with Enterprise Manager.

Configuring Per-Recipient Policies for DLP

You configure outgoing mail policies to use your DLP policies differently depending on whether you are using RSA Email DLP or RSA Enterprise Manager. For RSA Email DLP, you assign DLP policies to the mail policies using the Email Security appliance. For RSA Enterprise Manager, you assign the Email Security appliance's mail policies to DLP policies using Enterprise Manager.

RSA Email DLP

You enable RSA Email DLP policies on a per-recipient basis using the **Mail Policies > Outgoing Mail Policies** page (GUI) or the `policyconfig` command (CLI). You can enable different DLP policies for the different outgoing mail policies. You can only use DLP policies in outgoing mail policies.

DLP scanning takes place after the Outbreak Filters stage of the email "work queue." See Chapter 6, "Email Security Manager" for more information.

Figure 11-11 Default Outgoing Mail Policy with Enabled DLP Policies

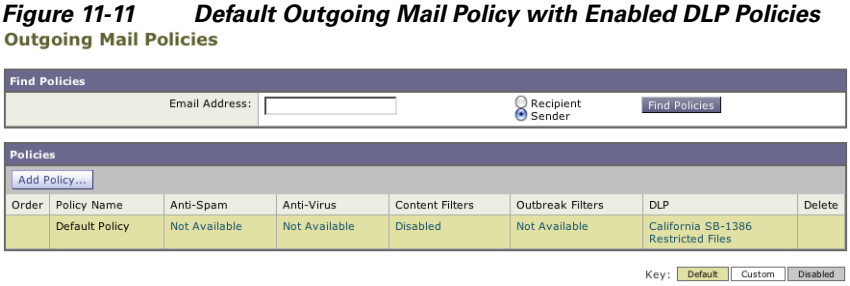

Editing the DLP Settings for a Mail Policy

The process for editing the per-user DLP settings for an outgoing mail policy is essentially the same for the default policy and individual policies. Individual policies (not the default) have an additional option for the DLP settings to **Enable DLP (Inherit default mail policy settings)**. Selecting this causes the policy to adopt all of the DLP settings from the default outgoing mail policy.

Figure 11-12 shows a list of DLP policies enabled for the default outgoing mail policy.

Figure 11-12 Enabling DLP Policies in the Default Outgoing Mail Policy

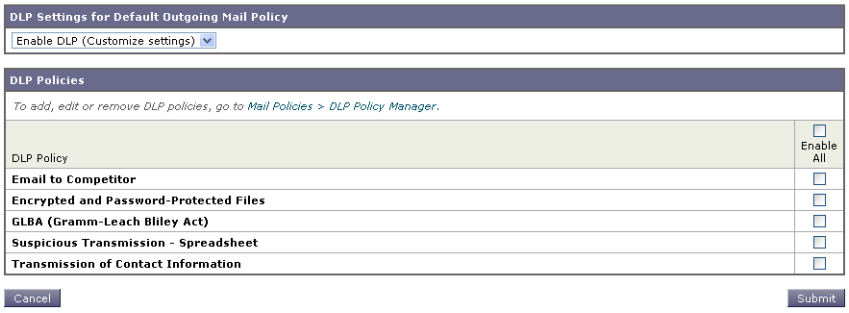

Step 1 Click the link for the DLP security service in any row of the Email Security Manager outgoing mail policy table. The DLP settings page is displayed.

Step 2 Click the link in the default row to edit the settings for the default policy.

Step 3 Select **Enable DLP (Customize Settings)** for the mail policy.

A list of the policies defined in the DLP Policy Manager is displayed.

Step 4 Select the RSA Email DLP policies that you want to use on this outgoing mail policy.

Step 5 Submit and commit your changes.

RSA Enterprise Manager

When using RSA Enterprise Manager for data loss prevention, you assign an outgoing mail policy to a DLP policy using Enterprise Manager instead of assigning a DLP policy to a mail policy using the Email Security appliance. See the RSA Enterprise Manager help for information on how to configure a DLP policy in Enterprise Manager.

Unlike with RSA Email DLP, outgoing mail policies cannot use the default policy's DLP policy when Enterprise Manager is enabled. If a mail policy is not specified for a DLP policy in Enterprise Manager, DLP scanning is not enabled on the mail policy.

If you try to delete an outgoing mail policy that is linked to a DLP policy, the Email Security appliance displays a message warning you that the mail policy is currently in use. If you decide to continue and delete the policy, Enterprise Manager automatically unlinks the deleted outgoing mail policy from any DLP policy that used it. Other than not scanning for messages based on the configuration of the deleted mail policy, DLP scanning will continue to work as before. The next DLP policy package sent to the Email Security appliance by Enterprise Manager will not include anything related to the deleted mail policy.

If you are managing multiple Email Security appliances with Enterprise Manager and you do not want one of the appliances to use a certain DLP policy, you can disable the DLP policies on an Email Security using the DLP Policy Manager. Go to Mail Policies > DLP Policy Manager and click **Enable or Disable Policies**. Clear the check box for any DLP policies you do not want to use on the Email Security appliance. See DLP Policy Manager for Enterprise Manager DLP Policies, page 11-30 for more information.

Figure 11-13 Enabling and Disabling DLP Policies Using DLP Policy Manager

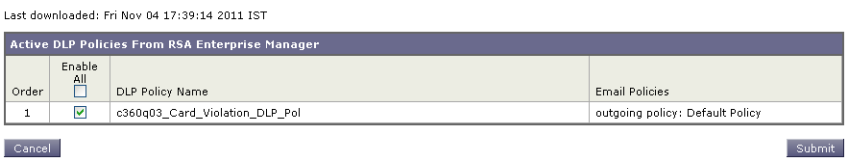

CHAPTER 12

Cisco IronPortEmail Encryption

Cisco IronPort AsyncOS supports using encryption to secure inbound and outbound email.

Cisco IronPortEmail Encryption: Overview

To use this feature, you create an encryption profile that specifies characteristics of the encrypted message and connectivity information for the key server. The key server may either be the Cisco Registered Envelope Service (managed service) or an Cisco IronPort Encryption appliance (locally managed server). Next, you create content filters or message filters (or both) to determine which messages to encrypt.

An outgoing message that meets the filter condition is placed in a queue on the Email Security appliance for encryption processing. Once the message is encrypted, the key used to encrypt it is stored on the key server specified in the encryption profile and the encrypted message is queued for delivery. If a temporary condition exists that prohibits the encryption of emails in the queue (i.e., temporary C-Series busyness or CRES unavailability), messages are re-queued and retried at a later time.

Note You can also set up the appliance to first attempt to send a message over a TLS connection before encrypting it. For more information, see Using a TLS Connection as an Alternative to Encryption, page 12-8.

Step 1 **If you want to use a local key server, configure the Cisco IronPort Encryption appliance.** For instructions on configuring key servers, see the *IronPort Encryption Appliance Local Key Server User Guide.*

Step 2 **Configure an encryption profile.** For instructions on configuring the encryption profile, see Configuring the Email Encryption Profile, page 12-3.

Step 3 **If you want to use the hosted key service, create a Cisco Registered Envelope Service corporate account. You create the account by clicking the Provision button after configuring an encryption profile.**

Step 4 **Configure an outgoing content filter.** You need to configure a content filter to tag the outbound emails that should be encrypted. For instructions on creating the content filter, see Configuring the Encryption Content Filter, page 12-7.

The following web browsers are supported:

- Microsoft® Internet Explorer 7 (Windows XP and Vista)
- Microsoft® Internet Explorer 8 (Windows XP and Vista)
- Firefox 3.0 and 3.5
- Safari 4.0 (Mac OS X)

Encryption Workflow

When using email encryption, the Cisco IronPort Email Security appliance encrypts a message and stores the message key on a local key server or a hosted key service. When the recipient opens an encrypted message, the recipient is authenticated by the key service, and the decrypted message is displayed.

Figure 12-1 Encryption Workflow

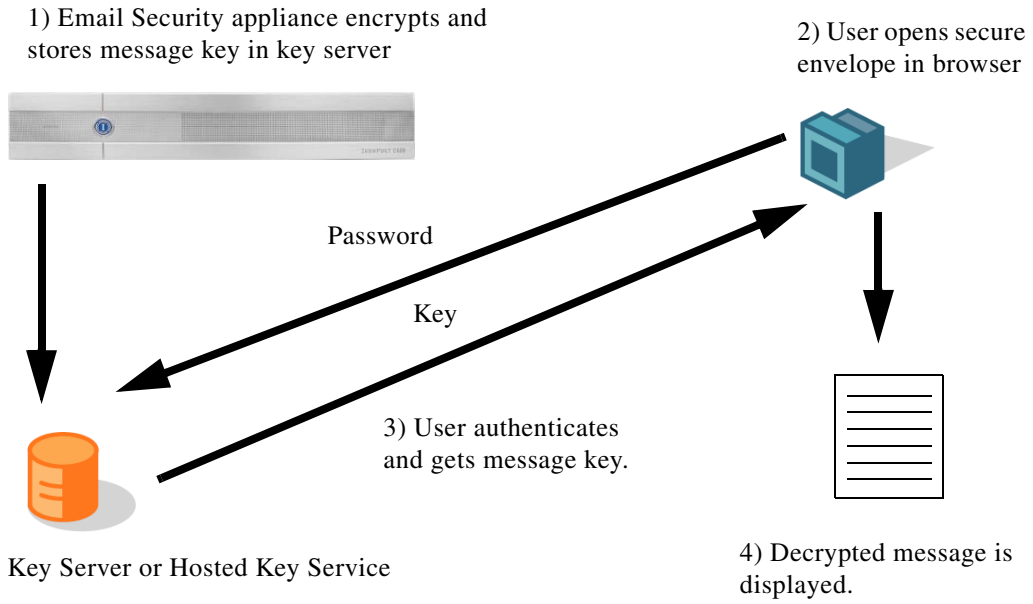

The basic workflow for opening encrypted messages is:

Step 1 When you configure an encryption profile, you specify the parameters for message encryption. For an encrypted message, the Email Security appliance creates and stores a message key on a local key server or on the hosted key service (Cisco Registered Envelope Service).

Step 2 The recipient opens the secure envelope in a browser.

Step 3 When a recipient opens an encrypted message in a browser, a password may be required to authenticate the recipient's identity. The key server returns the encryption key associated with the message.

Note When opening an encrypted email message for the first time, the recipient is required to register with the key service to open the secure envelope. After registering, the recipient may be able to open encrypted messages without authenticating, depending on settings configured in the encryption profile. The encryption profile may specify that a password isn't required, but certain features will be unavailable.

Step 4 The decrypted message is displayed.

Configuring the Email Encryption Profile

To use encryption with the Email Security appliance, you must configure an encryption profile. You can enable and configure an encryption profile using the `encryptionconfig` CLI command, or via Security Services > IronPort Email Encryption in the GUI.

Editing Email Encryption Global Settings

Step 1 Click Security Services > IronPort Email Encryption.

Step 2 Click **Enable**.

Step 3 Optionally, click **Edit Settings** and configure a proxy server.

Figure 12-2 Configuring Global Settings

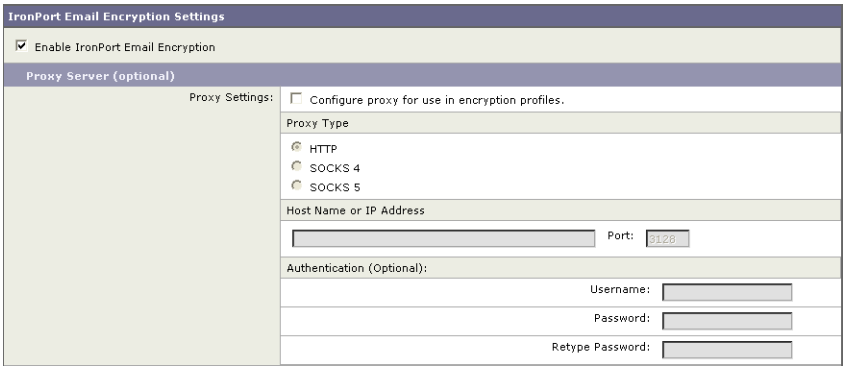

Adding an Encryption Profile

You can create one or more encryption profiles if you use a local key service. You might want to create different encryption profiles if you want to use different levels of security for different groups of email. For example, you might want messages containing sensitive material to be sent with high security, but other messages to be sent with medium security. In this case, you might create a high security encryption profile to associate with the messages containing certain key words (such as 'confidential'), and create another encryption profile for other outgoing messages.

You can assign an encryption profile to a custom user role to allow delegated administrators assigned to that role to use the encryption profile with their DLP policies and content filters. Only administrators, operators, and delegated users can use encryption profiles when configuring DLP policies and content

filters. Encryption profiles that are not assigned to a custom role are available for use by all delegated administrators with mail or DLP policy privileges. See the "Common Administrative Tasks" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide* for more information on

Note You can configure multiple encryption profiles for a hosted key service. If your organization has multiple brands, this allows you to reference different logos stored on the key server for the PXE envelopes.

You create and save an encryption profile to store the following encryption settings:

- **Key server settings**. Specify a key server and information for connecting to that key server.
- **Envelope settings**. Specify details about the message envelope, such as the level of security, whether to return read receipts, the length of time a message is queued for encryption before it times out, the type of encryption algorithm to use, and whether to enable a decryption applet to run on the browser.
- **Message settings**. Specify details about messages, such as whether to enable secure message forwarding and secure Reply All.
- **Notification settings**. Specify the notification template to use for text and HTML notifications, as well as encryption failure notifications. You create the templates in text resources and select the templates when creating the encryption profile. You can also specify a message subject for encryption failure notifications. For more information about notifications, see Encryption Notification Templates, page 14-30 and Bounce and Encryption Failure Notification Templates, page 14-27.

Figure 12-3 Adding an Encryption Envelope Profile

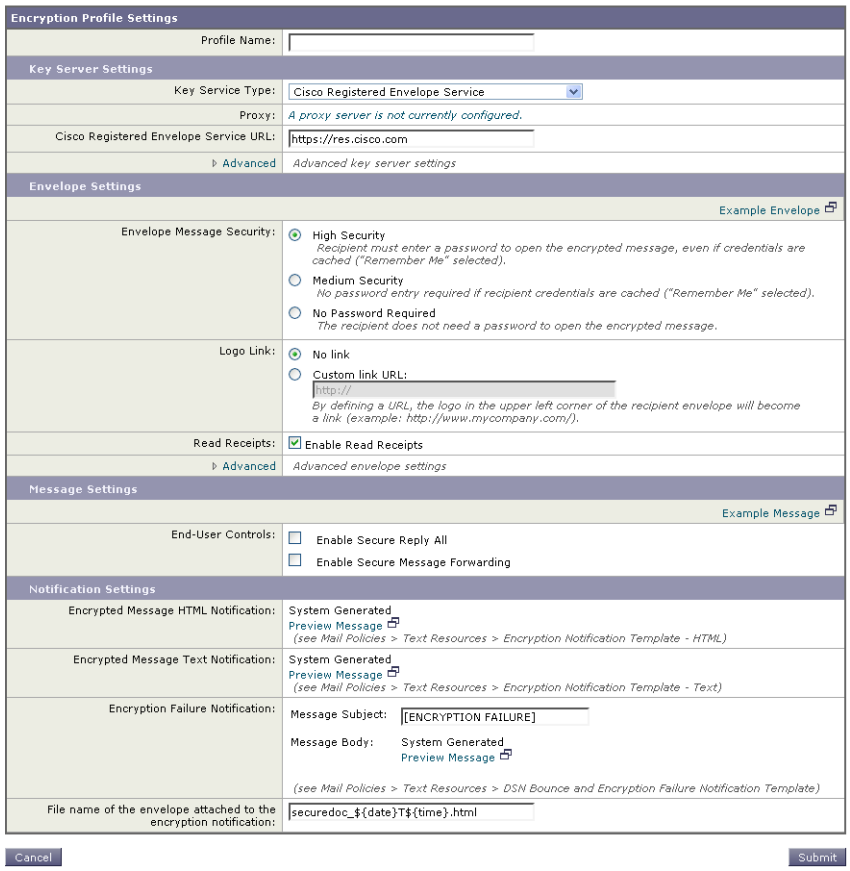

Step 1 In the Email Encryption Profiles section, click **Add Encryption Profile**.

Step 2 Enter a name for the Encryption Profile.

Step 3 Click the **Used By (Roles)** link, select the custom user role you want to have access to the encryption profile, and click **OK**.

Delegated administrators assigned to this custom role can use the encryption profile for any DLP policies and content filters for which they are responsible.

Step 4 In the Key Server Settings section, select from the following key servers:

- Cisco IronPort Encryption appliance (in network)
- Cisco Registered Envelope Service (hosted key service)

Step 5 If you select the Cisco IronPort Encryption appliance (local key service), enter the following settings:

- **Internal URL.** This URL is used by the Cisco IronPort Email Security appliance to contact the in-network Cisco IronPort Encryption appliance.
- **External URL**. This URL is used when the recipient's message accesses keys and other services on the Cisco IronPort Encryption appliance. The recipient uses this URL to make inbound HTTP or HTTPS requests.

Step 6 If you select the Cisco Registered Envelope Service, enter the URL for the hosted key service. The key service URL is `https://res.cisco.com`.

Step 7 Click **Advanced** under Key Server Settings to specify whether to use HTTP or HTTPS for transfering the envelope's encrypted payload when the recipient opens the envelope. You can from one of the following:

- **Use the Key Service with HTTP**. Transfers the encrpyted payload from the key service using HTTP when the recipient opens the envelope. If you are using Cisco Registered Envelope Service, this is the URL you specified in Step 6. If you are using the Cisco IronPort Encryption appliance, this is the external URL you specified in Step 5.

  Since the payload is already encrypted, transporting it over HTTP is safe and faster than sending over HTTPS. This provides better performance than sending image requests over HTTPS.
- **Use the Key Service with HTTPS**. Transfers the encrpyted payload from the key service using HTTPS when the recipient opens the envelope. If you are using Cisco Registered Envelope Service, this is the URL you specified in Step 6. If you are using the Cisco IronPort Encryption appliance, this is the external URL you specified in Step 5.
- **Specify a separate URL for payload transport**. If you don't want to use the key server for your encrypted payload, you can use another URL and specify whether to use HTTP or HTTPS for the payload transfer.

Step 8 In the Envelope Settings section, select the level of message security:

- **High Security**. The recipient must always enter a password to open encrypted messages.
- **Medium Security**. The recipient does not need to enter credentials to open the encrypted message if the recipient credentials are cached.
- **No Password Required**. This is the lowest level of encrypted message security. The recipient does not need to enter a password to open the encrypted message, but the read receipts, Secure Reply, Secure Reply All, and Secure Message Forwarding features will be unavailable to prevent another email user from sending a message on behalf of the original recipient.

Step 9 To enable users to open your organization's URL by clicking its logo, you can add a link to the logo. Choose from the following options:

- **No link.** A live link is not added to the message envelope.
- **Custom link URL**. Enter the URL to add a live link to the message envelope.

Step 10 Optionally, enable read receipts. If you enable this option, the sender receives a receipt when recipients open the secure envelope.

Step 11 Optionally, click **Advanced** under Envelope Settings to configure the following settings:

- Enter the length of time (in seconds) that a message can be in the encryption queue before timing out. Once a message times out, the appliance bounces the message and sends a notification to the sender.
- Select an encryption algorithm:
  - **ARC4**. ARC4 is the most common choice, providing strong encryption with minimal decryption delays for message recipients.
  - **AES**. AES provides stronger encryption but also takes longer to decrypt, introducing delays for recipients. AES is typically used in government and banking applications.
- Enable or disable the decryption applet. Enabling this option causes the message attachment to be opened in the browser environment. Disabling this option causes message attachments to be decrypted at the key server. If you disable this option, messages may take longer to open, but are not dependent on the browser environment.

Step 12 In the Message Settings section, enable or disable **Secure Reply All**.

Step 13 Enable or disable **Secure Message Forwarding**.

Step 14 Select an HTML notification template. Choose from HTML notifications you configured in text resources. If you did not configure a template, the system uses the default template.

Note The key server uses an HTML or text notification based on the recipient's email application. You must configure notifications for both.

Step 15 Select a text notification template. Choose from text notifications you configured in text resources. If you did not configure a template, the system uses the default template.

Step 16 Enter a subject header for encryption failure notifications. The appliance sends a notification if the encryption process times out.

Step 17 Select an encryption failure notification template for the message body. Choose from an encryption failure notification template you configured in text resources. If you did not configure a template, the system uses the default template.

Step 18 Submit and commit your changes.

Step 19 If you use Cisco Registered Envelope Service, you must take the additional step of provisioning your appliance. Provisioning the appliance registers the encryption profile with the hosted key service. To provision the appliance, click the **Provision** button for the encryption profile you want to register.

Updating the PXE Engine

The Cisco IronPort Email Encryption Settings page displays the current versions of the PXE engine and the Domain Mappings file used by your appliance. In previous versions of AsyncOS, you had to update AsyncOS in order to update the PXE engine. Now, you can use the Security Services > Service Updates page (or the `updateconfig` command in the CLI) to configure the Cisco IronPort appliance to automatically update the PXE engine. For more information, see Service Updates, page 15-10.

You can also manually update the engine using the **Update Now** button of the PXE Engine Updates section of IronPort Email Encryption Settings page (or the `encryptionupdate` command in the CLI).

Figure 12-4 PXE Engine Updates on the IronPort Email Encryption Settings Page

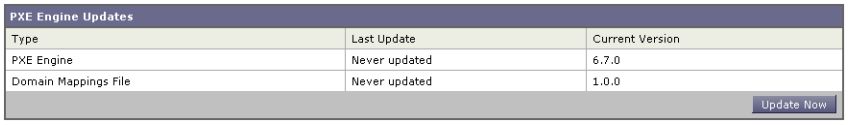

Configuring the Encryption Content Filter

After you create an encryption profile, you need to create an outgoing content filter that determines which email messages should be encrypted. The content filter scans outgoing email and determines if the message matches the conditions specified. Once the content filter determines a message matches the condition, the Cisco IronPort Email Security appliance encrypts the message and sends the generated key to the key server. It uses settings specified in the encryption profile to determine the key server to use and other encryption settings.

Using a TLS Connection as an Alternative to Encryption

Based on the destination controls specified for a domain, your Cisco IronPort appliance can securely relay a message over a TLS connection instead of encrypting it, if a TLS connection is available. The appliance decides whether to encrypt the message or send it over a TLS connection based on the TLS setting in the destination controls (Required, Preferred, or None) and the action defined in the encryption content filter.

When creating the content filter, you can specify whether to always encrypt a message or to attempt to send it over a TLS connection first, and if a TLS connection is unavailable, to encrypt the message. Table 12-1 shows you how an Email Security appliance will send a message based on the TLS settings for a domain's destination controls, if the encryption control filter attempts to send the message over a TLS connection first.

Table 12-1 TLS Support on ESA Appliances

Destination Controls TLS Setting	Action if TLS Connection Available	Action if TLS Connection Unavailable
None	Encrypt envelope and send	Encrypt envelope and send
TLS Preferred	Send over TLS	Encrypt envelope and send
TLS Required	Send over TLS	Retry/bounce message

For more information on enabling TLS on destination controls, see the "Customizing Listeners" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Creating a Content Filter to Encrypt and Deliver Now

Step 1 Go to Mail Policies > Outgoing Content Filters.

Step 2 In the Filters section, click **Add Filter**.

Step 3 In the Conditions section, click **Add Condition.**

Step 4 Add a condition to filter the messages that you want to encrypt. For example, to encrypt sensitive material, you might add a condition that identifies messages containing particular words or phrases, such as "Confidential," in the subject or body.

Step 5 Click **OK**.

For more details about building conditions, see Content Filters Overview, page 6-6.

Step 6 Optionally, click **Add Action** and select **Add Header** to insert an encryption header into the messages to specify an additional encryption setting.

For more information about encryption headers, see Inserting Encryption Headers into Messages, page 12-11.

Step 7 In the Actions section, click **Add Action**.

Step 8 Select **Encrypt and Deliver Now (Final Action)**.

Figure 12-5 Configuring the Encrypt and Deliver Now Action

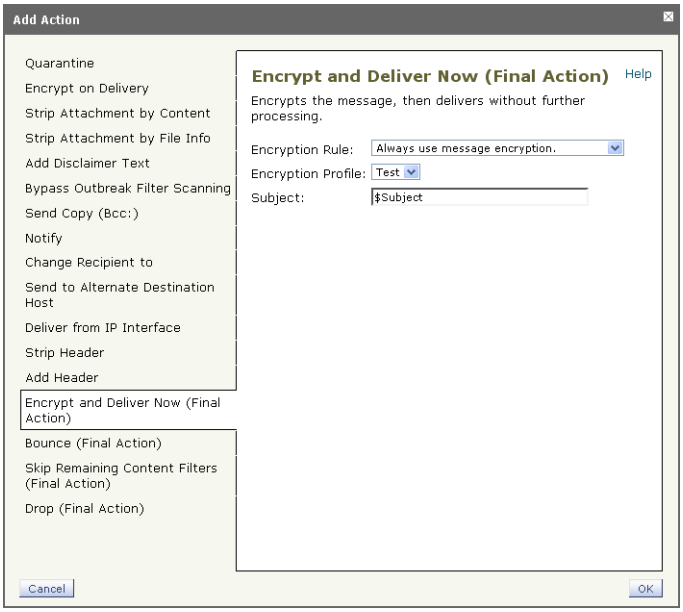

Step 9 Select whether to always encrypt messages that meet the condition or to only encrypt messages if the attempt to send it over a TLS connection fails.

Step 10 Select the encryption profile to associate with the content filter.

The encryption profile specifies settings about the key server to use, levels of security, formatting of the message envelope, and other message settings. When you associate an encryption profile with the content filter, the content filter uses these stored settings to encrypt messages.

Step 11 Enter a subject for the message.

Step 12 Click **OK**.

The content filter in Figure 12-6 shows a content filter that searches for ABA content in the message body. The action defined for the content filter specifies that the email is encrypted and delivered.

Figure 12-6 Encryption Content Filter

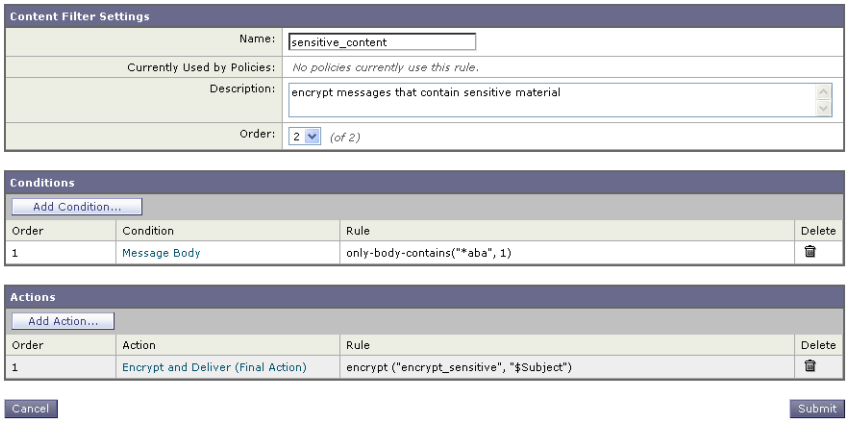

Step 13 After you add the encrypt action, click **Submit**.

Step 14 Commit your changes.

Step 15 Once you add the content filter, you need to add the filter to an outgoing mail policy. You may want to enable the content filter on the default policy, or you may choose to apply the filter to a specific mail policy, depending on your organization's needs. For information about working with mail policies, see Overview of User-Based Policies, page 6-1.

Creating a Content Filter to Encrypt on Delivery

To create a content filter to encrypt a message on delivery, which means that the message continues to the next stage of processing, and when all processing is complete, the message is encrypted and delivered:

Step 1 Go to Mail Policies > Outgoing Content Filters.

Step 2 In the Filters section, click **Add Filter**.

Step 3 In the Conditions section, click **Add Condition.**

Step 4 Add a condition to filter the messages that you want to encrypt. For example, to encrypt sensitive material, you might add a condition that identifies messages containing particular words or phrases, such as "Confidential," in the subject or body.

Step 5 Click **OK**.

For more details about building conditions, see Content Filters Overview, page 6-6.

Step 6 Optionally, click **Add Action** and select **Add Header** to insert an encryption header into the messages to specify an additional encryption setting.

For more information about encryption headers, see Inserting Encryption Headers into Messages, page 12-11.

Step 7 In the Actions section, click **Add Action**.

Step 8 Select **Encrypt on Delivery**.

Figure 12-7 Configuring the Encrypt on Delivery Action

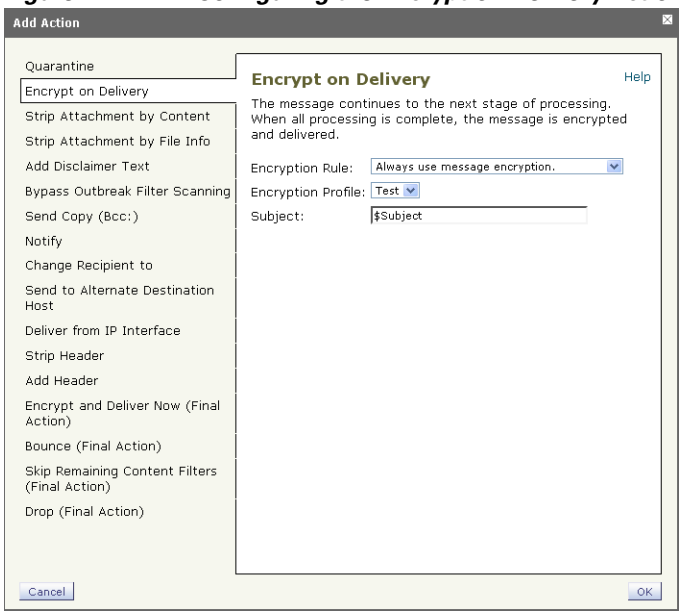

Step 9 Select whether to always encrypt messages that meet the condition or to only encrypt messages if the attempt to send it over a TLS connection fails.

Step 10 Select the encryption profile to associate with the content filter.

The encryption profile specifies settings about the key server to use, levels of security, formatting of the message envelope, and other message settings. When you associate an encryption profile with the content filter, the content filter uses these stored settings to encrypt messages.

Step 11 Enter a subject for the message.

Step 12 Click **OK**.

Step 13 After you add the encrypt action, click **Submit**.

Step 14 Commit your changes.

Step 15 Once you add the content filter, you need to add the filter to an outgoing mail policy. You may want to enable the content filter on the default policy, or you may choose to apply the filter to a specific mail policy, depending on your organization's needs. For information about working with mail policies, see Overview of User-Based Policies, page 6-1.

Inserting Encryption Headers into Messages

AsyncOS enables you to add encryption settings to a message by inserting an SMTP header into a message using either a content filter or a message filter. The encryption header can override the encryption settings defined in the associated encryption profile, and it can apply specified encryption features to messages.

To add an encryption header to a message by using a content filter, add the Add Header filter action to the content filter, and enter the encryption header and its value. For example, if you want a Registered Envelope to expire in 24 hours after you send it, type `X-PostX-ExpirationDate` as the header name and `+24:00:00` as the header value.

Figure 12-8 Configuring the Add Header Action

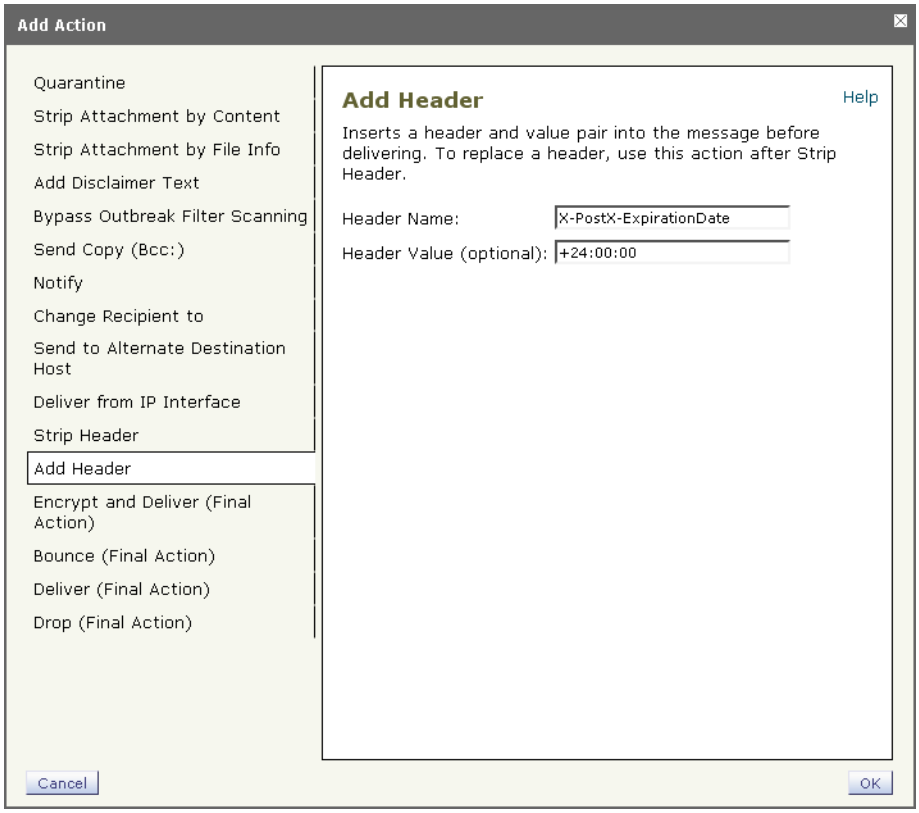

For more information about creating an encryption content filter, see Creating a Content Filter to Encrypt and Deliver Now, page 12-8. For information about inserting a header using a message filter, see the "Using Message Filters to Enforce Email Policies" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Encryption Headers

Table 12-2 displays the encryption headers that you can add to messages.

Table 12-2 Email Encryption Headers

MIME Header	Description	Value
`X-PostX-Reply-Enabled`	Indicates whether to enable secure reply for the message and displays the Reply button in the message bar. This header adds an encryption setting to the message.	A Boolean for whether to display the Reply button. Set to `true` to display the button. The default value is `false`.
`X-PostX-Reply-All-Enabled`	Indicates whether to enable secure "reply all" for the message and displays the Reply All button in the message bar. This header overrides the default profile setting.	A Boolean for whether to display Reply All button. Set to `true` to display the button. The default value is `false`.

Table 12-2 Email Encryption Headers

MIME Header	Description	Value
`X-PostX-Forward-Enabled`	Indicates whether to enable secure message forwarding and displays the Forward button in the message bar. This header overrides the default profile setting.	A Boolean for whether to display the Forward button. Set to `true` to display the button. The default value is `false`.
`X-PostX-Send-Return-Receipt`	Indicates whether to enable read receipts. The sender receives a receipt when recipients open the Secure Envelope. This header overrides the default profile setting.	A Boolean for whether to send a read receipt. Set to `true` to display the button. The default value is `false`.
`X-PostX-ExpirationDate`	Defines a Registered Envelope's expiration date before sending it. The key server restricts access to the Registered Envelope after the expiration date. The Registered Envelope displays a message indicating that the message has expired. This header adds an encryption setting to the message. If you use Cisco Registered Envelope Service, you can log in to the website at `http://res.cisco.com` and use the message management features to set, adjust, or eliminate the expiration dates of messages after you send them.	A string value containing relative date or time. Use the `+HH:MM:SS` format for relative hours, minutes, and seconds, and the `+D` format for relative days. By default, there is no expiration date.
`X-PostX-ReadNotificationDate`	Defines the Registered Envelope's "read by" date before sending it. The local key server generates a notification if the Registered Envelope has not been read by this date. Registered Envelopes with this header do not work with Cisco Registered Envelope Service, only a local key server. This header adds an encryption setting to the message.	A string value containing relative date or time. Use the `+HH:MM:SS` format for relative hours, minutes, and seconds, and the `+D` format for relative days. By default, there is no expiration date.
`X-PostX-Suppress-Applet-For-Open`	Indicates whether to disable the decryption applet. The decryption applet causes message attachments to be opened in the browser environment. Disabling the applet causes the message attachment to be decrypted at the key server. If you disable this option, messages may take longer to open, but they are not dependent on the browser environment. This header overrides the default profile setting.	A Boolean for whether to disable the decryption applet. Set to `true` to disable the applet. The default value is `false`.

Table 12-2 Email Encryption Headers

MIME Header	Description	Value
`X-PostX-Use-Script`	Indicates whether to send JavaScript-free envelopes. A JavaScript-free envelope is a Registered Envelope that does not include the JavaScript that is used to open envelopes locally on the recipient's computer. The recipient must use either the Open Online method or the Open by Forwarding method to view the message. Use this header if a recipient domain's gateway strips JavaScript and makes the encrypted message unopenable.This header adds an encryption setting to the message.	A Boolean for whether the JavaScript applet should be included or not. Set to `false` to send a JavaScript-free envelope. The default value is `true`.
`X-PostX-Remember-Envelope-Key-Checkbox`	Indicates whether to allow envelope-specific key caching for offline opening of envelopes. With envelope key caching, the decryption key for a particular envelope is cached on the recipient's computer when the recipient enters the correct password and selects the "Remember the password for this envelope" check box. After that, the recipient does not need to enter a password again to reopen the envelope on the computer. This header adds an encryption setting to the message.	A Boolean for whether to enable envelope key caching and display the "Remember the password for this envelope" check box. The default value is `false`.

Encryption Headers Examples

This section provides examples of encryption headers.

Enabling Envelope Key Caching for Offline Opening

To send a Registered Envelope with envelope key caching enabled, insert the following header into the message:

```
X-PostX-Remember-Envelope-Key-Checkbox: true
```

The "Remember the password for this envelope" check box is displayed on the Registered Envelope.

Enabling JavaScript-Free Envelopes

To send a Registered Envelope that is JavaScript-free, insert the following header into the message:

```
X-PostX-Use-Script: false
```

When the recipient opens the securedoc.html attachment, the Registered Envelope is displayed with an Open Online link, and the Open button is disabled.

Enabling Message Expiration

To configure a message so that it expires 24 hours after you send it, insert the following header into the message:

```
X-PostX-ExpirationDate: +24:00:00
```

The recipient can open and view the content of the encrypted message during the 24-hour period after you send it. After that, the Registered Envelope displays a message indicating that the envelope has expired.

Disabling the Decryption Applet

To disable the decryption applet and have the message attachment decrypted at the key server, insert the following header into the message:

X-PostX-Suppress-Applet-For-Open: true

Note The message may take longer to open when you disable the decryption applet, but it is not dependent on the browser environment.

CHAPTER 13

SenderBase Network Participation

SenderBase is an email reputation service designed to help email administrators research senders, identify legitimate sources of email, and block spammers.

In the System Setup Wizard (GUI) and the `systemsetup` command (CLI) you can agree to participate in the SenderBase Network. Cisco will collect aggregated email traffic statistics about your organization. This includes only summary data on message attributes and information on how different types of messages were handled by Cisco IronPort appliances. For example, Cisco does not collect the message body or the message subject. Personally identifiable information or information that identifies your organization will be kept confidential.

Sharing Statistics with SenderBase

Step 1 Access the Security Services > SenderBase page.

Figure 13-1 Security Services > SenderBase Page

Note If you have not already agreed to the license agreement during system setup (see Step 2: System, page 3-15), this page will look different. You must click **Enable** on the Security Services > SenderBase page and then read and agree to the license before you can edit global settings.

Step 2 Click **Edit Global Settings**.

Figure 13-2 Security Services > SenderBase page: Edit

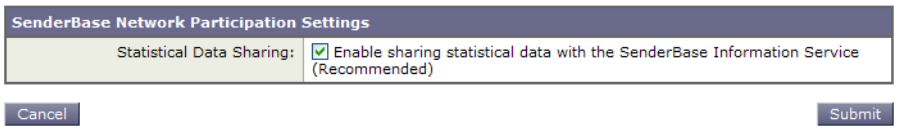

Step 3 Mark the box to enable sharing statistical data with the SenderBase Information Service. Checking this box enables the feature globally for the appliance. When enabled, the Context Adaptive Scanning Engine (CASE) is used to collect and report the data (regardless of whether or not Cisco IronPort anti-spam scanning is enabled).

Step 4 As an option, you can enable a proxy server for sharing statistical data with the SenderBase Information Service. If you define a proxy server to retrieve rules updates, you can also configure an authenticated username, password, and specific port when connecting to the proxy server in the additional fields provided. To edit these settings, see System Time, page 15-47. You can configure the same settings using the `senderbaseconfig` command in the CLI

Frequently Asked Questions

Cisco recognizes that privacy is important to you, so we design and operate our services with the protection of your privacy in mind. If you enroll in SenderBase Network Participation, Cisco will collect aggregated statistics about your organization's email traffic; however, we do not collect or use any personally identifiably information. Any information Cisco collects that would identify your users or your organization will be treated as confidential.

Why should I participate?

Participating in the SenderBase Network helps us help you. Sharing data with us is important to helping stop email-based threats such as spam, viruses and directory harvest attacks from impacting your organization. Examples of when your participation is especially important include:

- Email attacks that are specifically targeted at your organization, in which case the data you contribute provides the primary source of information to protect you.
- Your organization is one of the first to be hit by a new global email attack, in which case the data you share with us will dramatically improve the speed with which we are able to react to a new threat.

What data do I share?

The data is summarized information on message attributes and information on how different types of messages were handled by Cisco IronPort appliances. We do not collect the full body of the message. Again, information provided to Cisco that would identify your users or your organization will be treated as confidential. (See What does Cisco do to make sure that the data I share is secure?, page 13-4 below).

Table 13-1 and Table 13-2 explain a sample log entry in a "human-friendly" format.

Table 13-1 Statistics Shared Per Cisco IronPort Appliance

Item	Sample Data
MGA Identifier	MGA 10012
Timestamp	Data from 8 AM to 8:05 AM on July 1, 2005
Software Version Numbers	MGA Version 4.7.0
Rule Set Version Numbers	Anti-Spam Rule Set 102
Anti-virus Update Interval	Updates every 10 minutes
Quarantine Size	500 MB
Quarantine Message Count	50 messages currently in quarantine
Virus Score Threshold	Send messages to quarantine at threat level 3 or higher
Sum of Virus Scores for messages entering quarantine	120
Count of messages entering quarantine	30 (yields average score of 4)
Maximum quarantine time	12 hours
Count of Outbreak quarantine messages broken down by why they entered and exited quarantine, correlated with Anti-Virus result	50 entering quarantine due to .exe rule 30 leaving quarantine due to manual release, and all 30 were virus positive
Count of Outbreak quarantine messages broken down by what action was taken upon leaving quarantine	10 messages had attachments stripped after leaving quarantine
Sum of time messages were held in quarantine	20 hours

Table 13-2 Statistics Shared Per IP Address

Item	Sample Data
Message count at various stages within the appliance	Seen by Anti-Virus engine: 100 Seen by Anti-Spam engine: 80
Sum of Anti-Spam and Anti-Virus scores and verdicts	2,000 (sum of anti-spam scores for all messages seen)
Number of messages hitting different Anti-Spam and Anti-Virus rule combinations	100 messages hit rules A and B 50 messages hit rule A only
Number of Connections	20 SMTP Connections
Number of Total and Invalid Recipients	50 total recipients 10 invalid recipients
Hashed Filename(s): (a)	A file <one-way-hash>.pif was found inside an archive attachment called <one-way-hash>.zip.

Table 13-2 Statistics Shared Per IP Address

Item	Sample Data (Continued)
Obfuscated Filename(s): (b)	A file aaaaaaa0.aaa.pif was found inside a file aaaaaaa.zip.
URL Hostname (c)	There was a link found inside a message to www.domain.com
Obfuscated URL Path (d)	There was a link found inside a message to hostname www.domain.com, and had path aaa000aa/aa00aaa.
Number of Messages by Spam and Virus Scanning Results	10 Spam Positive 10 Spam Negative 5 Spam Suspect 4 Virus Positive 16 Virus Negative 5 Virus Unscannable
Number of messages by different Anti-Spam and Anti-Virus verdicts	500 spam, 300 ham
Count of Messages in Size Ranges	125 in 30K-35K range
Count of different extension types	300 ".exe" attachments
Correlation of attachment types, true file type, and container type	100 attachments that have a ".doc" extension but are actually ".exe" 50 attachments are ".exe" extensions within a zip
Correlation of extension and true file type with attachment size	30 attachments were ".exe" within the 50-55K range

(a) Filenames will be encoded in a 1-way hash (MD5).

(b) Filenames will be sent in an obfuscated form, with all lowercase ASCII letters ([a-z]) replaced with "a," all uppercase ASCII letters ([A-Z]) replaced with "A," any multi-byte UTF-8 characters replaced with "x" (to provide privacy for other character sets), all ASCII digits ([0-9]) replaced with "0," and all other single byte characters (whitespace, punctuation, etc.) maintained. For example, the file Britney1.txt.pif would appear as Aaaaaaa0.aaa.pif.

(c) URL hostnames point to a web server providing content, much as an IP address does. No confidential information, such as usernames and passwords, are included.

(d) URL information following the hostname is obfuscated to ensure that any personal information of the user is not revealed.

What does Cisco do to make sure that the data I share is secure?

If you agree to participate in the SenderBase Network:

Data sent from your Cisco IronPort appliances will be sent to the Cisco IronPort SenderBase Network servers using the secure protocol HTTPS.

All customer data will be handled with care at Cisco. This data will be stored in a secure location and access to the data will be limited to employees and contractors at Cisco who require access in order to improve the company's email security products and services or provide customer support.

No information identifying email recipients or the customer's company will be shared outside of Cisco Systems when reports or statistics are generated based on the data.

Will sharing data impact the performance of my Cisco IronPort appliances?

Cisco believes that there will be a minimal performance impact for most customers. We record data that already exists as part of the mail delivery process. Customer data is then aggregated on the appliance and sent to SenderBase servers in batches, typically every 5 minutes. We anticipate that the total size of data transferred via HTTPS will be less than 1% of the bandwidth of a typical company's email traffic.

When enabled, the Context Adaptive Scanning Engine (CASE) is used to collect and report the data (regardless of whether or not Cisco IronPort anti-spam scanning is enabled).

Note If you choose to participate in the SenderBase Network, a "body scan" is performed on each message. This happens regardless of whether a filter or other action applied to the message would have triggered a body scan. See "Body Scanning Rule" in the "Using Message Filters to Enforce Email Policies" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information about body scanning.

If you have additional questions, please contact Cisco IronPort Customer Support. See Cisco IronPort Support Community, page 1-10.

Are there other ways I can share data?

For customers that want to share additional data to do even more to help Cisco provide top quality security services, there is also a command that allows this. This higher level of data sharing will also provide attachment filenames in clear, unhashed text, as well as hostnames of URLs in messages. If you are interested in learning more about this feature, please talk to your Systems Engineer or contact Cisco IronPort Customer Support.

CHAPTER 14

Text Resources

Overview

This chapter discusses creating and managing various text resources, such as content dictionaries, DLP dictionaries, disclaimers, and templates.

Content Dictionaries

You can use content dictionaries to scan messages against message or content filters in order to take appropriate action in accordance with your corporate policies. You can create, delete, and view dictionaries; add and delete entries from a dictionary; and import and export entire dictionaries. You can also determine case sensitivity and word boundary detection for each dictionary. For example, you could create a list of confidential or profane words, and, using a filter rule to scan messages for words in the list, drop or archive messages containing matching words. And you can add a "weight" terms in a dictionary so that certain terms trigger a filter action more easily.

Dictionaries can contain non-ASCII characters.

DLP Dictionaries

You can use data los prevention (DLP) dictionaries in custom DLP policies to scan outgoing messages for sensitive information. Similar to content dictionaries, you can create, delete, and view dictionaries; add and delete entries from a dictionary; and import and export entire dictionaries. Unlike content dictionaries, terms in DLP policies do not have a "weight." AsyncOS comes with a set of predefined dictionaries from RSA Security Inc., but you can create custom DLP dictionaries.

Dictionary terms are case-sensitive and can contain non-ASCII characters. For more information on data loss prevention, see Chapter 11, "Data Loss Prevention."

Text Resources

Text resources are text objects, such as disclaimers, notification templates, and anti-virus templates. You can create new objects for use in various components of AsyncOS. You can import and export text resources.

Message Disclaimer Stamping

Message disclaimer stamping allows you to add a disclaimer text resource to messages. For example, you could append a copyright statement, promotional message, or disclaimer to every message sent from within your enterprise.

Content Dictionaries

AsyncOS provides two types of dictionaries: content and DLP dictionaries. For information on managing DLP dictionaries, see DLP Dictionaries, page 14-9.

Content dictionaries are groups of words or entries that work in conjunction with the Body Scanning feature on the appliance and are available to both content and message filters. Use the dictionaries you define to scan messages, message headers, and message attachments for terms included in the dictionary in order to take appropriate action in accordance with your corporate policies. For example, you could create a list of confidential or profane words, and, using a filter rule to scan messages that contain words in the list, drop, archive, or quarantine the message.

The AsyncOS operating system includes the ability to define a total of 100 content dictionaries using the GUI (Mail Policies > Dictionaries) or the CLI's `dictionaryconfig` command. You can create, delete, and view dictionaries; add and delete entries from a dictionary; and import and export entire dictionaries.

Dictionary Content

Words in dictionaries are created with one text string per line, and entries can be in plain text or in the form of regular expressions. Dictionaries can also contain non-ASCII characters. Defining dictionaries of regular expressions can provide more flexibility in matching terms, but doing so requires you to understand how to delimit words properly. For a more detailed discussion of Python style regular expressions, consult the Python Regular Expression HOWTO, accessible from

```
http://www.python.org/doc/howto/
```

Note To use the special character # at the beginning of a dictionary entry, you can use a character class [#] to prevent it being treated as a comment.

For each term, you specify a "weight," so that certain terms can trigger filter conditions more easily. When AsyncOS scans messages for the content dictionary terms, it "scores" the message by multiplying the number of term instances by the weight of term. Two instances of a term with a weight of three would result in a score of six. AsyncOS then compares this score with a threshold value associated with the content or message filter to determine if the message should trigger the filter action.

You can also add smart identifiers to a content dictionary. Smart identifiers are algorithms that search for patterns in data that correspond to common numeric patterns, such as social security numbers and ABA routing numbers. These identifiers can useful for policy enforcement. For more information about regular expressions, see "Regular Expressions in Rules" in the "Using Message Filters to Enforce Email Policies" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*. For more information about smart identifiers, see "Smart Identifiers" in the "Using Message Filters to Enforce Email Policies" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Note Dictionaries containing non-ASCII characters may or may not display properly in the CLI on your terminal. The best way to view and change dictionaries that contain non-ASCII characters is to export the dictionary to a text file, edit that text file, and then import the new file back into the appliance. For more information, see Importing and Exporting Dictionaries as Text Files, page 14-3.

Word Boundaries and Double-byte Character Sets

In some languages (double-byte character sets), the concepts of a word or word boundary, or case do not exist. Complex regular expressions that depend on concepts like what is or is not a character that would compose a word (represented as "\w" in regex syntax) cause problems when the locale is unknown or if the encoding is not known for certain. For that reason, you may want to disable word-boundary enforcement.

Importing and Exporting Dictionaries as Text Files

The content dictionary feature also includes, by default, the following text files located in the configuration directory of the appliance:

- `config.dtd`
- `profanity.txt`
- `proprietary_content.txt`
- `sexual_content.txt`

These text files are intended to be used in conjunction with the content dictionaries feature to aid you in creating new dictionaries. These content dictionaries are weighted and use smart identifiers to better detect patterns in data and trigger filters when the patterns indicate compliance issues.

Note Importing and exporting dictionaries does not preserve the Match Whole Words and Case Sensitive settings. This settings are only preserved in the configuration file.

See Appendix A, "Accessing the Appliance" for more information accessing on the configuration directory.

You can also create your own dictionary files and import them onto the appliance. The best way to add non-ASCII characters to dictionaries is to add the terms into the dictionary in a text file off the appliance, move that file onto the appliance, and then import that file as a new dictionary. For more information about importing dictionaries, see Importing Dictionaries, page 14-6. For information about exporting dictionaries, see Exporting Dictionaries, page 14-7.

You can also import and export custom DLP dictionaries. For more information, see Importing and Exporting DLP Dictionaries, page 14-11.

Warning **These text files contain terms that some persons may consider obscene, indecent or offensive. If you import terms from these files into your content dictionaries, the terms will be displayed when you later view the content dictionaries you have configured on the appliance.**

Managing Content Dictionaries (GUI)

Log in to the GUI and click the Mail Policies tab. Click the Dictionaries link in the left menu.

Figure 14-1 The Dictionaries Page

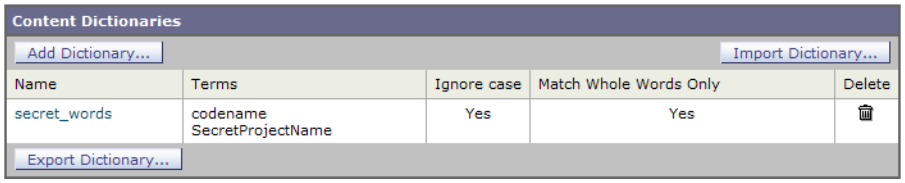

Adding Dictionaries

Step 1 Click **Add Dictionary** on the Dictionaries page. The Add Dictionary page is displayed:

Figure 14-2 The Dictionaries Page

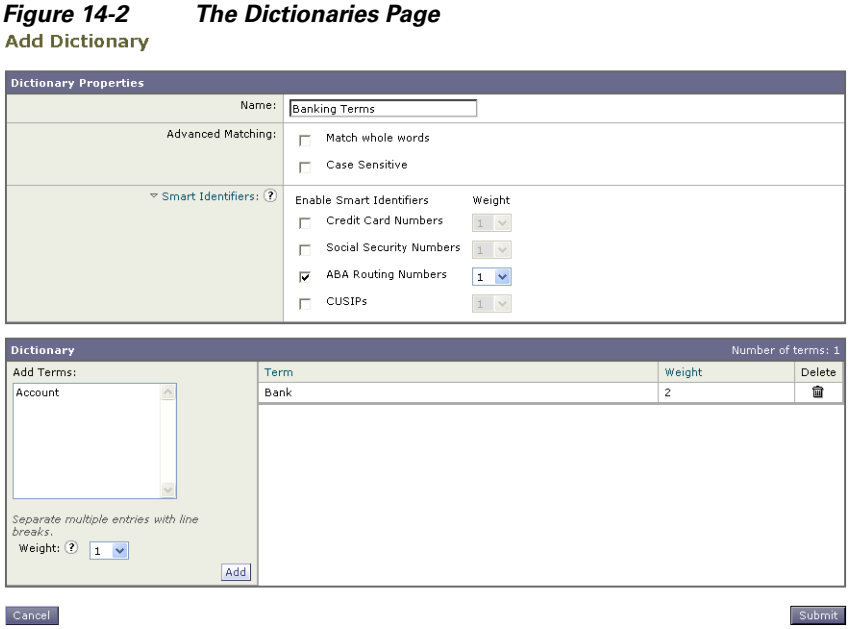

Step 2 Type a name for the dictionary.

Step 3 Specify whether to match whole words only by marking the checkbox next to Match Whole Words Only. See Matching Whole Words Only, page 14-6 for more information.

Step 4 Specify whether to perform case-sensitive searches. See Matching Case-Sensitive Words, page 14-6 for more information.

Note AysncOS preserves the Match Whole Words and Case Sensitive settings when they are saved in the configuration file. These settings are not preserved when importing and exporting dictionaries.

Step 5 Optionally, add a smart-identifier to the dictionary. Smart identifiers are algorithms that search for patterns in data that correspond to common numeric patterns, such as social security numbers and ABA routing numbers. For more information about smart identifiers, see the "Using Message Filters to Enforce Email Policies" chapter in *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Step 6 Enter new dictionary entries into the list of terms. For more information about the kinds of entries that are supported, see Dictionary Content, page 14-2.

Step 7 Specify a weight for the term. You can "weight" a dictionary term so that it is more likely than other terms to trigger a filter action. For more information about how this weight is used to determine filter actions, see "Threshold Scoring for Content Dictionaries" in the "Using Message Filters to Enforce Email Policies" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Step 8 Click **Add**.

Step 9 Submit and commit your changes.

The Dictionaries page now lists the new dictionary, along with the terms included and the setting configured for the dictionary.

Note Content dictionary entries with the regular expression: ". *" at the beginning or end will cause the system to lock if a match for the "word" MIME part is found. Cisco Systems recommends you do not use ". *" at the beginning or end of a content dictionary entry.

Matching Case-Sensitive Words

Checking this box will cause AsyncOS to consider the case of the word when matching. For example, the words “codename” would match a dictionary entry of “codename”, but the word “CodeName” would not match.

Matching Whole Words Only

Checking this box will cause words to match only if they match the whole entry. For example, the word “codename” would match a dictionary entry of “codename,” while “code” and “codenam” would not.

Sorting Terms

You can click the column heading to sort by term or weight. If you click the column heading a second time, it reverses the sort order.

Editing Dictionaries

Step 1 Click the name of the dictionary in the listing on the Dictionaries page. The Edit Dictionary page is displayed.

Step 2 Make changes to the entries or the settings for the dictionary, and click **Submit**.

Step 3 Commit your changes.

Deleting Dictionaries

Step 1 Click the trash can icon next to the dictionary to delete in the dictionary listing. A confirmation message is displayed.

Step 2 The confirmation message lists any filters that are currently referencing the dictionary.

Step 3 Click **Delete** to delete the dictionary.

Step 4 Commit your changes.

Step 5 Any message filters that reference the deleted dictionary are marked as invalid.

Step 6 Any content filters that reference the deleted dictionary are left enabled, but will evaluate to false.

Importing Dictionaries

Step 1 Click **Import Dictionary** on the Dictionaries page. The Import Dictionary dialog is displayed:

Figure 14-3 The Import Dictionary Page

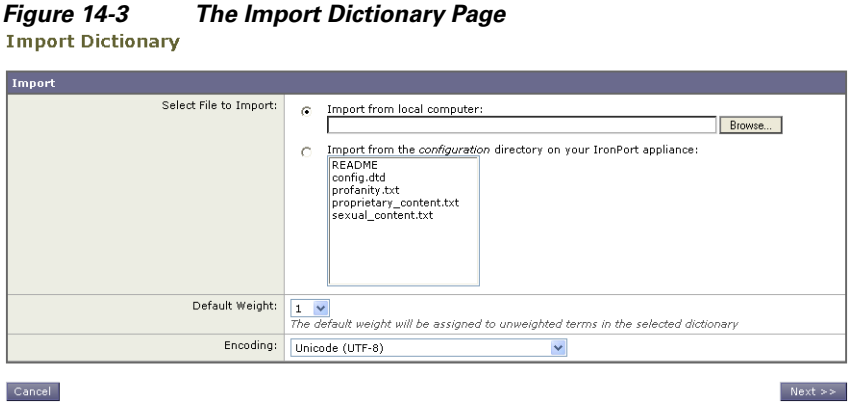

Step 2 Select the location to import from.

Step 3 Select a file to import.

Note The file to import must be in the configuration directory on the appliance.

Step 4 Select the default weight to use for dictionary terms. AsyncOS will assign a default weight to any terms with unspecified weights. You can edit the weights after importing the file.

Step 5 Select an encoding.

Step 6 Click **Next**.

Step 7 The imported dictionary is displayed in the Add Dictionary page.

Step 8 You can now name and edit the dictionary before adding it.

Step 9 Submit and commit your changes.

Exporting Dictionaries

Step 1 Click **Export Dictionary** on the Dictionaries page. The Export Dictionary dialog is displayed:

Figure 14-4 The Export Dictionary Page

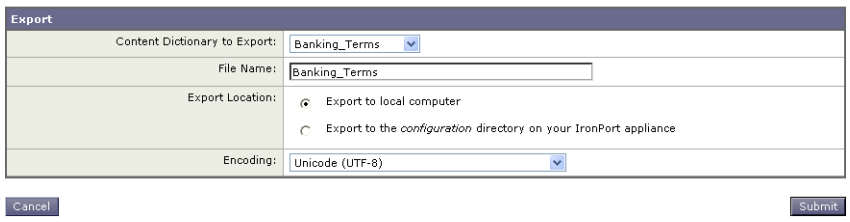

Step 2 Select a dictionary to export.

Step 3 Enter a file name for the dictionary. This is the name of the file that will be created in the configuration directory on the appliance.

Step 4 Select the location to export to.

Step 5 Select an encoding for the text file.

Step 6 Submit and commit your changes.

Using and Testing Content Dictionaries

Dictionaries can be used along with the various `dictionary-match()` message filter rules and with content filters.

Dictionary Match Filter Rule

The message filter rule named `dictionary-match(<dictionary_name>)` (and its counterparts) evaluates to true if the message body contains any of the regular expressions in the content dictionary named *dictionary_name*. If that dictionary does not exist, the rule evaluates to false.

Note that the `dictionary-match()` rule functions similarly to the `body-contains()` body scanning rule: it only scans the body and attachments of messages, and not the headers.

For scanning headers, you can use the appropriate `*-dictionary-match()`-type rule (there are rules for specific headers, such as `subject-dictionary-match()` and a more generic rule, `header-dictionary-match()`, in which you can specify any header including custom headers). See "Dictionary Rules" in the "Using Message Filters to Enforce Email Policies" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information about dictionary matching.

Table 14-1 Message Filter Rules for Content Dictionaries

Rule	Syntax	Description
Dictionary Match	`dictionary-match(<dictionary_name>)`	Does the message contain a word that matches all the regular expressions listed in the named dictionary?

In the following example, a new message filter using the `dictionary-match()` rule is created to blind carbon copy the administrator when the Cisco IronPort appliance scans a message that contains any words within the dictionary named "secret_words" (created in the previous example). Note that because of the settings, only messages that contain the whole word "`codename`" matching the case exactly will evaluate to true for this filter.

```
bcc_codenames:

   if (dictionary-match ('secret_words'))

       {

       bcc('administrator@example.com');

       }
```

In this example, we send the message to the Policy quarantine:

```
quarantine_codenames:

   if (dictionary-match ('secret_words'))

       {
```

```
quarantine('Policy');

}
```

Example Dictionary Entries

Table 14-2 Example Dictionary Entries

Description	Example
Wildcard	*
Anchors	Ends with: foo$ Begins with: ^foo
Email address (Do not escape the period)	**foo@example.com**, **@example.com** **example.com$** (ends with) **@example.***
Subject	**An email subject** (keep in mind when using the ^ anchor in email subjects that subjects are often prepended with “RE:” or “FW:” and the like)

Testing Content Dictionaries

The `trace` function can provide quick feedback on message filters that use the `dictionary-match()` rule. See Debugging Mail Flow Using Test Messages: Trace, page -446 for more information. You can also use the `quarantine()` action to test filters, as in the `quarantine_codenames` filter example above.

DLP Dictionaries

DLP dictionaries are groups of words or phrases that work in conjunction with the RSA DLP scanning feature on the appliance and are available to custom DLP policies. Use the DLP dictionaries to scan messages and message attachments for the words and phrases included in the dictionary in order to take appropriate action in accordance with your corporate policies. AsyncOS comes with a set of predefined dictionaries from RSA Security Inc., but you can create custom DLP dictionaries.

You can also create your own dictionary as a text file on your local machine and import it onto the appliance. Use line breaks for each term in the dictionary text file. Dictionary terms are case-sensitive and can contain non-ASCII characters.

You manage DLP dictionaries using the DLP Policy Manager. To open the DLP Policy Manager, select the Mail Policies > DLP Policy Manager menu in the GUI. For more information on the DLP Policy Manager, see Chapter 11, “Data Loss Prevention.”

Adding Custom Dictionaries

Step 1 Click the **Custom DLP Dictionaries** link in the DLP Policy Manager.

The DLP Dictionaries page appears.

Step 2 Click Add Dictionary.

The Add Dictionary Page appears.

Figure 14-5 Add DLP Dictionaries

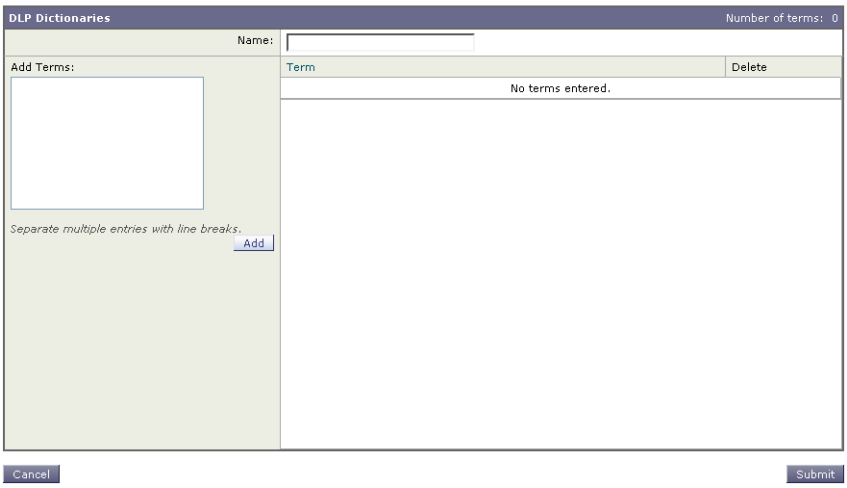

Step 3 Enter a name for the custom dictionary.

Step 4 Enter new dictionary entries into the list of terms. You can use line breaks to enter multiple entries at once.

Step 5 Click **Add**.

Step 6 Submit and commit the new dictionary.

The Dictionaries page now lists the new dictionary, along with the terms included and the setting configured for the dictionary.

Editing Custom DLP Dictionaries

Step 1 Click on the name of the dictionary in the listing on the DLP Dictionaries page.

Step 2 Make changes to the entries.

Step 3 Submit and commit your changes.

Deleting Custom DLP Dictionaries

Step 1 Click the trash can icon next to the dictionary to delete in the dictionary listing. A confirmation message is displayed listing any filters that are currently referencing the dictionary.

Step 2 Click **Delete** to delete the dictionary.

Step 3 Commit your changes.

Importing and Exporting DLP Dictionaries

You can create your own DLP dictionary as a text file on your local machine and import it into AsyncOS, as well as export existing custom dictionaries as text files. Predefined DLP dictionaries cannot be exported.

The DLP dictionary file includes a list of the words and phrases used as dictionary terms with line breaks separating each term. If you export an existing content dictionary to use as a DLP dictionary, you need to strip the weight values from the text file and convert any regular expressions to words or phrases before importing the file as a DLP dictionary.

Importing DLP Dictionaries as a Text File

Step 1 Click Import Dictionary on the DLP Dictionaries page.

The Import Dictionary dialog is displayed:

Figure 14-6 Importing Dictionaries

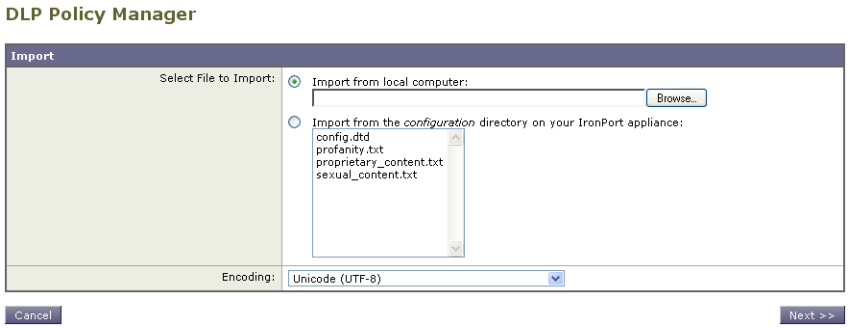

Step 2 Select a file to import from either your local machine or the configuration directory on the appliance.

Step 3 Select an encoding.

Step 4 Click **Next**.

Step 5 The imported dictionary is displayed in the Add Dictionary page.

Step 6 You can now name and edit the dictionary before adding it.

Step 7 Submit and commit your changes.

Exporting DLP Dictionaries as a Text File

Step 1 Click Export Dictionary on the Dictionaries page.

The Export Dictionary dialog is displayed:

Figure 14-7 Exporting Dictionaries

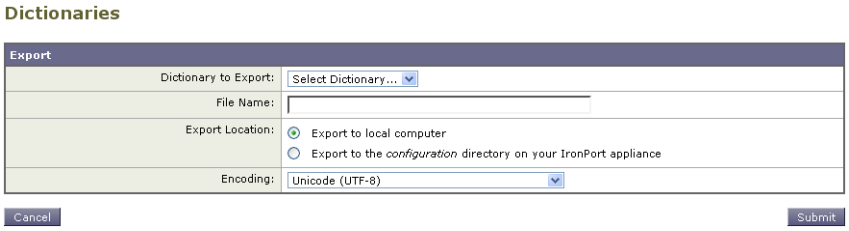

Step 2 Select a dictionary to export.

Step 3 Enter a file name for the dictionary.

Step 4 Choose where to save the exported dictionary, either on your local computer or in the configuration directory on the appliance.

Step 5 Select an encoding for the file.

Step 6 Submit and commit your changes.

Understanding Text Resources

Text resources are text templates that can be attached to messages or sent as messages. Text resources can be one of the following types:

- **Message disclaimers —** Text that is added to messages. For more information, see Disclaimer Template, page 14-17.
- **Notification templates —** Messages that are sent as notifications, used with the `notify()` and `notify-bcc()` actions. For more information, see Notification Templates, page 14-24.
- **Anti-virus Notification templates —** Messages that are sent as notifications when a virus is found in a message. You can create a template for a container (which appends the original message), or as a notice that is sent without the appended message. For more information, see Anti-Virus Notification Templates, page 14-24.
- **Bounce and Encryption Failure Notification templates —** Messages that are sent as notifications when a message is bounced or message encryption fails. For more information, see Bounce and Encryption Failure Notification Templates, page 14-27.
- **DLP Notification templates —** Messages that are sent when an email message contains information that violates your organization's data loss prevention policies. For more information, see DLP Notification Templates, page 14-28.
- **Encryption Notification Templates —** Messages that are sent when you configure the Cisco IronPort appliance to encrypt outgoing email. The message notifies recipients that they have received an encrypted message and provides instructions for reading it. For more information, see Encryption Notification Templates, page 14-30.

You can use the CLI (`textconfig`) or the GUI to manage text resources, including: adding, deleting, editing, importing, and exporting. For information on managing text resources using the GUI, see Managing Text Resources (GUI), page 14-13.

Text resources can contain non-ASCII characters.

Note Text resources containing non-ASCII characters may or may not display properly in the CLI on your terminal. To view and change text resources that contain non-ASCII characters, export the text resource to a text file, edit that text file, and then import the new file back into the appliance. For more information, see Importing and Exporting Text Resources as Text Files, page 14-13.

Importing and Exporting Text Resources as Text Files

You must have access to the configuration directory on the appliance. Imported text files must be present in the configuration directory on the appliance. Exported text files are placed in the configuration directory.

See Appendix A, "Accessing the Appliance" for more information on accessing the configuration directory.

To add non-ASCII characters to text resources, add the terms into the text resource in a text file off the appliance, move that file onto the appliance, and then import that file as a new text resource. For more information about importing text resources, see Importing Text Resources, page 14-14. For information about exporting text resources, see Exporting Text Resources, page 14-15.

Managing Text Resources (GUI)

You can manage text resources in the GUI on the Mail Policies > Text Resources page. From the Text Resources page you can add, edit, delete, export, and import text resources.

You can define plain text messages for all text resource types, and you can also define HTML-based messages for some text resource types. For more information, see Working with HTML-Based Text Resources, page 14-16.

Figure 14-8 The Text Resources Page

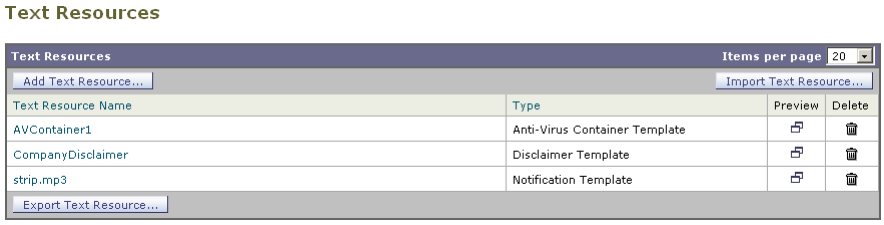

Note You can manage text resources from the CLI using the `textconfig` command.

Adding Text Resources

Step 1 On the Mail Policies > Text Resources page, click **Add Text Resource**. The Add Text Resource page is displayed.

Figure 14-9 Adding a Text Resource

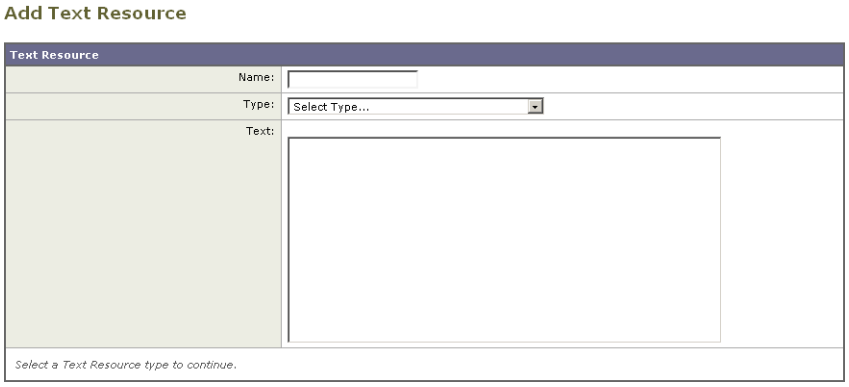

Step 2 Enter a name for the text resource in the Name field.

Step 3 Select the type of text resource from the Type field.

Step 4 Enter the message text in the appropriate field. If the text resource allows only plain text messages, use the Text field. If the text resource allows both HTML and plain text messages, use the HTML and Plain Text fields. For more information, see Working with HTML-Based Text Resources, page 14-16.

Step 5 Submit and commit your changes.

Editing Text Resources

Step 1 On the Mail Policies > Text Resources page, click the name of the text resource you want to edit. The Edit Text Resource page is displayed.

Step 2 Make changes to the text resource.

Step 3 Submit and commit your changes.

Deleting Text Resources

You can delete text resources from the Text Resources page. However, please note the following impact:

- Any message filters that reference the deleted text resource are marked as invalid.
- Any content filters that reference the deleted text resource are left enabled, but will evaluate to false.

Step 1 On the Mail Policies > Text Resources page, click the trash can icon under the Delete column for the text resource you want to delete. A confirmation message is displayed.

Step 2 Click **Delete** to delete the text resource.

Step 3 Commit your changes.

Importing Text Resources

Step 1 On the Mail Policies > Text Resources page, click **Import Text Resource**. The Import Text Resource page is displayed.

Figure 14-10 Importing a Text Resource

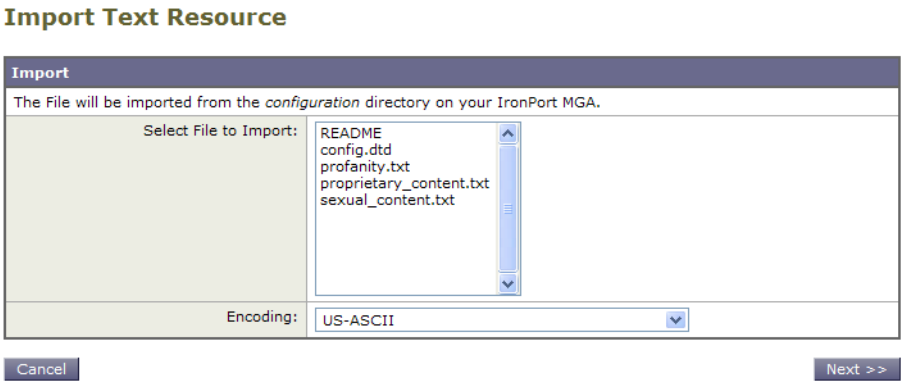

Step 2 Select a file to import.

Note The file to import must be in the configuration directory on the appliance.

Step 3 Specify an encoding.

Step 4 Click **Next**.

The imported text resource is displayed in the Text field of the Add Text Resource page.

Step 5 Choose a name, edit, and select the text resource type.

Step 6 Submit and commit your changes.

Exporting Text Resources

When you export a text resource, a text file is created in the configuration directory on the appliance.

Step 1 On the Mail Policies > Text Resources page, click **Export Text Resource**. The Export Text Resource dialog is displayed:

Figure 14-11 Exporting a Text Resource

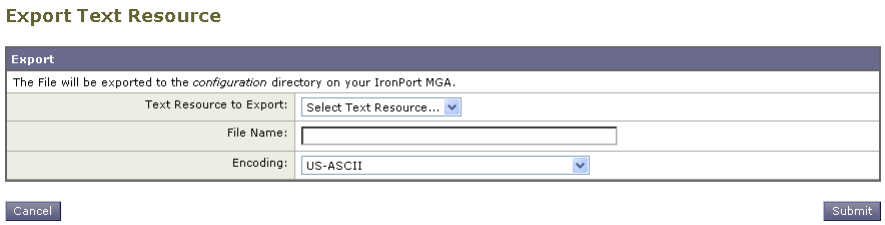

Step 2 Select a text resource to export.

Step 3 Enter a file name for the text resource.

Step 4 Select an encoding for the text file.

Step 5 Click **Submit** to create the text file containing the text resource in the configuration directory.

Working with HTML-Based Text Resources

You can create some text resources with both HTML-based and plain text messages, such as Disclaimers. When a text resource containing both HTML-based and plain text messages is applied to an email message, the HTML-based text resource message is applied to the text/html part of the email message, and the plain text message is applied to the text/plain part of the email message.

When you add or edit an HTML-based text resource, the GUI includes a rich text edit that allows you to enter rich text without having to manually write HTML code.

Figure 14-12 shows the rich text editor for an HTML-based text resource.

Figure 14-12 Creating HTML-Based Text Resources

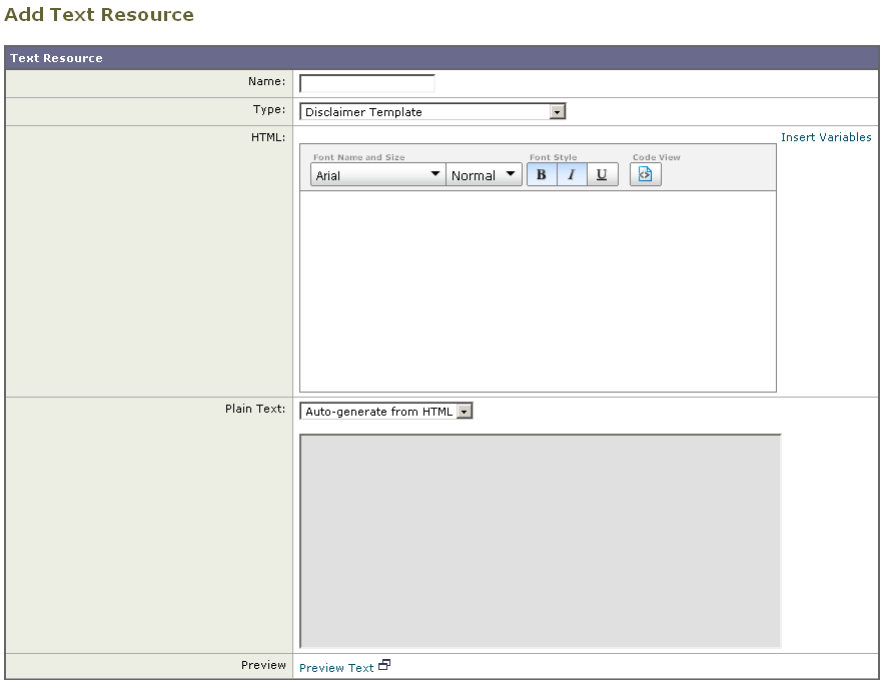

Consider the following rules and guidelines when adding and editing an HTML-based text resource:

- You can choose to have the plain text version of the message to be automatically generated based on the HTML version, or you can define the plain text version independently.
- You can switch between the rich text editor and HTML code by clicking the **Code View** button.
- To enter HTML code that is not supported in the rich text editor in the GUI, switch to code view and manually enter HTML code. For example, you might want to do this to insert a reference to an image file located on an external server using the `<img src>` HTML tag.

Importing and Exporting HTML-Based Text Resources

You can export to and import from a text file HTML-based text resources. When you export an HTML-based text resource to a file, the file contains the following sections for each version of the text resource:

```
[html_version]
[text_version]
```

The order of these sections does not matter.

For example, an exported file might contain the following text:

```
[html_version]
<p>Sample <i>message.</i></p>
[text_version]
Sample message.
```

Consider the following rules and guidelines when exporting and importing HTML-based text resources:

- When you export an HTML-based text resource whose plain text message is automatically generated from the HTML version, the exported file does not contain the `[text_version]` section.
- When you import from a text file, any HTML code under the `[html_version]` section is converted to the HTML message in the created text resource if the text resource type supports HTML messages. Similarly, any text under the `[text_version]` section is converted to the plain text message in the created text resource.
- When you import from a file that contains an empty or nonexistent `[html_version]` section to create a HTML-based text resource, the Cisco IronPort appliance creates both an HTML and plain text message using the text in the `[text_version]` section.

Using Text Resources

All types of text resources are created in the same way, using the Text Resources page or the `textconfig` CLI command. Once created, each type is used in a different way. Disclaimers and notification templates are used with filters and listeners, while anti-virus notification templates are used with mail policies and anti-virus settings.

Disclaimer Template

The Cisco IronPort appliance can add a default disclaimer above or below the text (heading or footer) for some or all messages received by a listener. You can add disclaimers to messages on the Cisco IronPort appliance using the following methods:

- Via a listener, using the GUI or the `listenerconfig` command (see Adding Disclaimer Text via a Listener, page 14-18).
- Using the content filter action, `Add Disclaimer Text` (see Content Filter Actions, page 6-12).
- Using the message filter action, `add-footer()` (see the "Using Message Filters to Enforce Email Policies" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*).
- Using a data loss prevention profile (see Data Loss Prevention, page 11-1).
- Using message modification for Outbreak Filters to alert the user that the message may be an attempt at phishing or malware distribution (see Modifying Messages, page 10-6). Disclaimers added for this type of notification are added above the text.

For example, you can append a copyright statement, promotional message, or disclaimer to every message sent from within your enterprise.

Prior to using disclaimer text you have to create the disclaimer template. Use the Text Resources page in the GUI (see Adding Text Resources, page 14-13) or the `textconfig` command (see the *Cisco IronPort AsyncOS CLI Reference Guide*) to create and manage a set of text strings to be used.

Adding Disclaimer Text via a Listener

Once you have disclaimer text resources created, select which text strings will be appended to messages received by the listener. You can add disclaimer text above or below a message. This feature is available on both public (inbound) and private (outbound) listeners.

If you send a message that consists of text and HTML (Microsoft Outlook calls this type of message a "multipart alternative"), the Cisco IronPort appliance will stamp the disclaimer on both parts of the message. However, if your message has signed content, the content will not be modified because the modification will invalidate the signature. Instead, a new part is created with a disclaimer stamp that says "`Content-Disposition inline attachment`." For more information on multipart messages, see "Message Bodies vs. Message Attachments" in the "Using Message Filters to Enforce Email Policies" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

The following example shows how to select a disclaimer to apply to messages on a listener via the GUI:

Figure 14-13 Editing a Listener to Include a Disclaimer

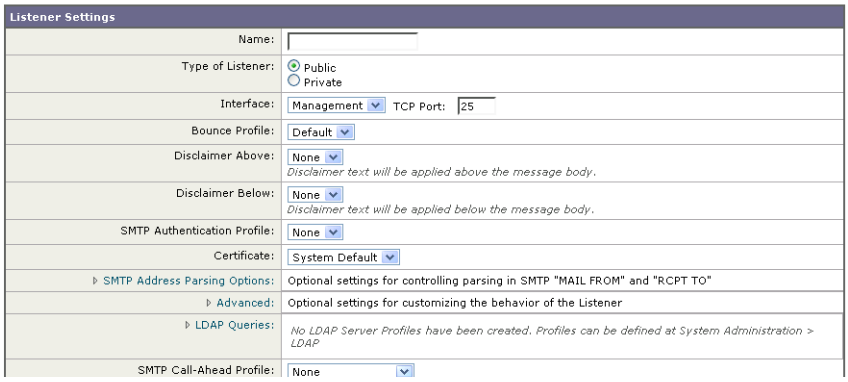

Adding Disclaimers via Filters

You can also append specific, predefined text strings to the disclaimers of messages using the filter action `add-footer()` or the content filter action "Add Disclaimer Text." For example, the following message filter rule appends the text string named `legal.disclaimer` to all messages sent from users in the LDAP group "Legal:"

```
Add-Disclaimer-For-Legal-Team:

if (mail-from-group == 'Legal')

{

  add-footer('legal.disclaimer');

}
```

Disclaimers and Filter Action Variables

You can also use message filter action variables (see "Action Variables" in the "Using Message Filters to Enforce Email Policies" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information).

The following variables are available for the Disclaimer Template:

Table 14-3 Anti-Virus Notification Variables

Variable	Substituted With
$To	Replaced by the message To: header (not the Envelope Recipient).
$From	Replaced by the message From: header (not the Envelope Sender).
$Subject	Replaced by the subject of the original message.
$Date	Replaced by the current date, using the format MM/DD/YYYY.
$Time	Replaced by the current time, in the local time zone.
$GMTimestamp	Replaced by the current time and date, as would be found in the Received: line of an email message, using GMT.
$MID	Replaced by the Message ID, or "MID" used internally to identify the message. Not to be confused with the RFC822 "Message-Id" value (use $Header to retrieve that).
$Group	Replaced by the name of the sender group the sender matched on when injecting the message. If the sender group had no name, the string ">Unknown<" is inserted.
$Policy	Replaced by the name of the HAT policy applied to the sender when injecting the message. If no predefined policy name was used, the string ">Unknown<" is inserted.
$Reputation	Replaced by the SenderBase Reputation score of the sender. If there is no reputation score, it is replaced with "None".
$filenames	Replaced with a comma-separated list of the message's attachments' filenames.
$filetypes	Replaced with a comma-separated list of the message's attachments' file types.
$filesizes	Replaced with a comma-separated list of the message's attachment's file sizes.
$remotehost	Replaced by the hostname of the system that sent the message to the Cisco IronPort appliance.
$AllHeaders	Replaced by the message headers.
$EnvelopeFrom	Replaced by the Envelope Sender (Envelope From, <MAIL FROM>) of the message.
$Hostname	Replaced by the hostname of the Cisco IronPort appliance.
$header['*string*']	Replaced by the value of the quoted header, if the original message contains a matching header. Note that double quotes may also be used.
$enveloperecipients	Replaced by all Envelope Recipients (Envelope To, <RCPT TO>) of the message.
$bodysize	Replaced by the size, in bytes, of the message.
$FilterName	Returns the name of the filter being processed.

Table 14-3 Anti-Virus Notification Variables (Continued)

Variable	Substituted With
$MatchedContent	Returns the content that triggered a scanning filter rule (including filter rules such as `body-contains` and content dictionaries).
$DLPPolicy	Replaced by the name of the email DLP policy violated.
$DLPSeverity	Replaced by the severity of violation. Can be “Low,” “Medium,” “High,” or “Critical.”
$DLPRiskFactor	Replaced by the risk factor of the message’s sensitive material (score 0 - 100).
$threat_category	Replaced with the type of Outbreak Filters threat, such as phishing, virus, scam, or malware.
$threat_type	Replaced by a subcategory of the Outbreak Filters threat category. For example, can be a charity scam, a financial phishing attempt, a fake deal, etc.
$threat_description	Replaced by a description of the Outbreak Filters threat.
$threat_level	Replaced by the message’s threat level (score 0 - 5).

To use message filter action variables in disclaimers, create a message disclaimer (via the Text Resource page in the GUI or the **`textconfig`** command), and reference the variable:

```
(running textconfig command)


Enter or paste the message disclaimer here. Enter '.' on a blank line to end.
This message processed at: $Timestamp
.


Message disclaimer "legal.disclaimervar" created.


Current Text Resources:
1. legal.disclaimer (Message Disclaimer)
2. legal.disclaimervar (Message Disclaimer)


Choose the operation you want to perform:
- NEW - Create a new text resource.
- IMPORT - Import a text resource from a file.
```

```
- EXPORT - Export text resource to a file.

- PRINT - Display the content of a resource.

- EDIT - Modify a resource.

- DELETE - Remove a resource from the system.

[]>


mail3.example.com>commit
```

Now, use the new disclaimer in a filter

```
Add-Timestamp:

if (mail-from-group == 'Legal')

{

  add-footer('legal.disclaimervar');

}
```

The `add-footer()` action supports non-ASCII text by adding the footer as an inline, UTF-8 coded, quoted printable attachment.

Disclaimer Stamping and Multiple Encodings

AsyncOS includes a setting used to modify the way disclaimer stamping with different character encodings works. By default, AsyncOS attempts to place the disclaimers it attaches within the body part of an email message. You can use a setting configured within the `localeconfig` command to configure the behavior if the encodings of the body part and the disclaimer are different. To understand this setting, it is helpful to view an email message as consisting of several parts:

To: joe@example.com From: mary@example.com Subject: Hi!	Headers
<blank line>	
Hello!	Body part
This message has been scanned...	First attachment part
Example.zip	Second attachment part

The message body after the first blank line may contain many MIME parts. The second and following parts are often called “attachments,” while the first is often called the “body” or “text.”

A disclaimer can be included in an email as either an attachment (above) or as part of the body

To: joe@example.com From: mary@example.com Subject: Hi!	Headers
<blank line>	
Hello!	Body part
This message has been scanned...	Disclaimer now included in body part
Example.zip	First attachment part

Typically, when there is an encoding mismatch between the message body and a disclaimer, AsyncOS attempts to encode the entire message in the same encoding as the message body so that the disclaimer will be included in the body (“inline”) and not included as a separate attachment. In other words, the disclaimer will be included inline if the encoding of the disclaimer matches that of the body, or if the text in the disclaimer contains characters that can be displayed inline (in the body). For example, it is possible to have a ISO-8859-1 encoded disclaimer that only contains US-ASCII characters; consequently, this will display “inline” without problems.

However, if the disclaimer cannot be combined with the body, you can use the `localeconfig` command to configure AsyncOS to attempt to promote, or convert, the body text to match the encoding of the disclaimer so that the disclaimer can be included in the body of the message:

```
example.com> localeconfig


Behavior when modifying headers: Use encoding of message body

Behavior for untagged non-ASCII headers: Impose encoding of message body

Behavior for mismatched footer or heading encoding: Only try encoding from

message body


Choose the operation you want to perform:

- SETUP - Configure multi-lingual settings.

[]> setup


If a header is modified, encode the new header in the same encoding as

the message body? (Some MUAs incorrectly handle headers encoded in a

different encoding than the body. However, encoding a modified header
```

```
in the same encoding as the message body may cause certain characters in the modified
header to be lost.) [Y]>


If a non-ASCII header is not properly tagged with a character set and

is being used or modified, impose the encoding of the body on the

header during processing and final representation of the message?

(Many MUAs create non-RFC-compliant headers that are then handled in

an undefined way. Some MUAs handle headers encoded in character sets

that differ from that of the main body in an incorrect way. Imposing the encoding of the
body on the header may encode

the header more precisely. This will be used to interpret the content of headers for
processing, it will not modify or rewrite the header

 unless that is done explicitly as part of the processing.) [Y]>


Footers or headings are added in-line with the message body whenever

possible. However, if the footer or heading is encoded differently

than the message body, and if imposing a single encoding will cause

loss of characters, it will be added as an attachment. The system will

always try to use the message body's encoding for the footer or

heading. If that fails, and if the message body's encoding is US-

ASCII, the system can try to edit the message body to use the footer's

or heading's encoding. Should the system try to impose the footer's

or headings's encoding on the message body? [N]> y


Behavior when modifying headers: Use encoding of message body

Behavior for untagged non-ASCII headers: Impose encoding of message

body. Behavior for mismatched footer or heading encoding: Try both

body and footer or heading encodings


Choose the operation you want to perform:
```

```
- SETUP - Configure multi-lingual settings.
```

For more information about the localeconfig command, see the "Customizing Listeners" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Notification Templates

Notification templates are used with the `notify()` and `notify-copy()` filter actions. Notification templates may contain non-ascii text and action variables (see "Action Variables" in the "Using Message Filters to Enforce Email Policies" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*), including the anti-virus-related variables used by anti-virus notifications. For example, you could use the `$Allheaders` action variable to include the headers from the original message. You can configure the From: address for notifications, see Configuring the Return Address for Various Generated Messages, page 15-15.

Once you have created a notification template, you can refer to it in content and message filters. Figure 14-14 shows a content filter where the `notify-copy()` filter action is set to send the "grape_text" notification to "grapewatchers@example.com:"

Figure 14-14 Notify Example in a Content Filter

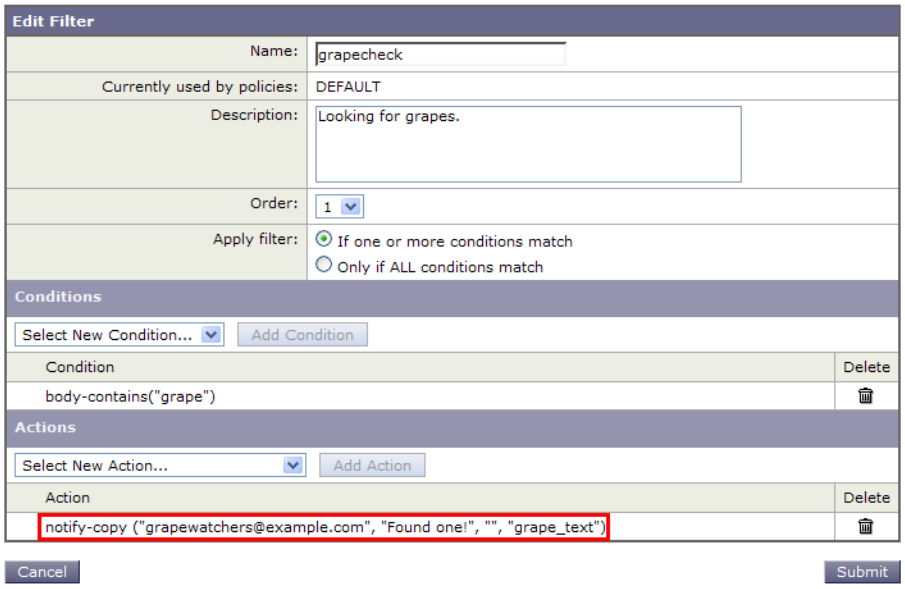

Anti-Virus Notification Templates

There are two types of anti-virus notification templates:

- **anti-virus notification template**. The anti-virus notification template is used when the original message is not attached to the virus notification.

- **anti-virus container template**. The container template is used when the original message is sent as an attachment.

Anti-virus notification templates are used in basically the same way as notification templates except that they are used with the anti-virus engine instead of filters. You can specify a custom notification to send while editing a mail policy. You can configure the From: address for anti-virus notifications. For information, see Configuring the Return Address for Various Generated Messages, page 15-15.

Custom Anti-Virus Notification Templates

Figure 14-15 shows a mail policy where a custom anti-virus notification is specified.

Figure 14-15 Anti-Virus Container Template Notification Example in a Mail Policy

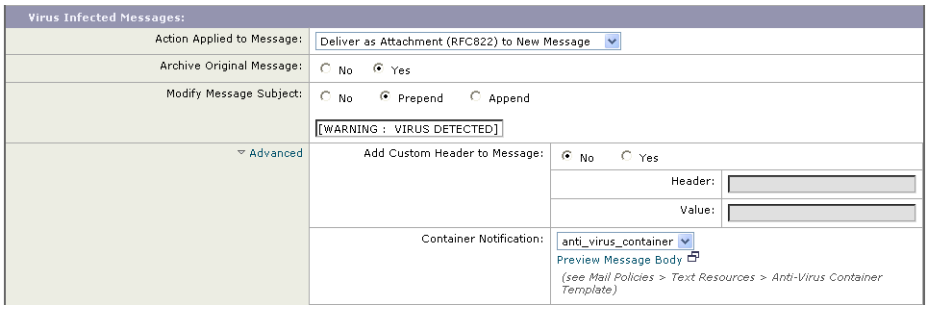

Anti-Virus Notification Variables

When creating an anti-virus notification, you can use any of the notification variables listed in Table 14-4:

Table 14-4 Anti-Virus Notification Variables

Variable	Substituted With
$To	Replaced by the message To: header (not the Envelope Recipient).
$From	Replaced by the message From: header (not the Envelope Sender).
$Subject	Replaced by the subject of the original message.
$AV_VIRUSES	Replaced by the list of all the viruses found anywhere in the message: “Unix/Apache.Trojan”, “W32/Bagel-F”
$AV_VIRUS_TABLE	Replaced by the table of MIME-Part/Attachment names and viruses in each part: “HELLO.SCR” : “W32/Bagel-F” <unnamed part of the message> : “Unix/Apache.Trojan”
$AV_VERDICT	Replaced by the anti-virus verdict.

Table 14-4 Anti-Virus Notification Variables (Continued)

Variable	Substituted With
$AV_DROPPED_TABLE	Replaced by the table of attachments that were dropped. Each row is composed of a part or filename followed by the list of viruses associated with that part: "HELLO.SCR" : "W32/Bagel-f", "W32/Bagel-d" "Love.SCR" : "Netsky-c", "W32/Bagel-d"
$AV_REPAIRED_VIRUSES	Replaced by the list of all the viruses found and repaired.
$AV_REPAIRED_TABLE	Replaced by the table of all parts and viruses found and repaired: "HELLO.SCR" : "W32/Bagel-F"
$AV_DROPPED_PARTS	Replaced by the list of filenames that were dropped: "HELLO.SCR", "CheckThisOut.exe"
$AV_REPAIRED_PARTS	Replaced by the list of filenames or parts that were repaired.
$AV_ENCRYPTED_PARTS	Replaced by the list of filenames or parts that were encrypted.
$AV_INFECTED_PARTS	Replaced by a comma-separated list of filenames for the files that contained a virus.
$AV_UNSCANNABLE_PARTS	Replaced by the list of filenames or parts that were unscannable.
$Date	Replaced by the current date, using the format MM/DD/YYYY.
$Time	Replaced by the current time, in the local time zone.
$GMTimestamp	Replaced by the current time and date, as would be found in the Received: line of an email message, using GMT.
$MID	Replaced by the Message ID, or "MID" used internally to identify the message. Not to be confused with the RFC822 "Message-Id" value (use $Header to retrieve that).
$Group	Replaced by the name of the sender group the sender matched on when injecting the message. If the sender group had no name, the string ">Unknown<" is inserted.
$Policy	Replaced by the name of the HAT policy applied to the sender when injecting the message. If no predefined policy name was used, the string ">Unknown<" is inserted.
$Reputation	Replaced by the SenderBase Reputation score of the sender. If there is no reputation score, it is replaced with "None".
$filenames	Replaced with a comma-separated list of the message's attachments' filenames.
$filetypes	Replaced with a comma-separated list of the message's attachments' file types.
$filesizes	Replaced with a comma-separated list of the message's attachment's file sizes.
$remotehost	Replaced by the hostname of the system that sent the message to the Cisco IronPort appliance.
$AllHeaders	Replaced by the message headers.

Table 14-4 Anti-Virus Notification Variables (Continued)

Variable	Substituted With
$EnvelopeFrom	Replaced by the Envelope Sender (Envelope From, <MAIL FROM>) of the message.
$Hostname	Replaced by the hostname of the Cisco IronPort appliance.

Note Variable names are not case-sensitive. For example, specifying "$to" is equivalent to specifying "$To" in the text resource. If an "AV_" variable is empty in the original message, the string `<None>` is substituted.

After the text resource has been defined, use the Mail Policies > Incoming/Outgoing Mail Policies > Edit Anti-Virus Settings page or the `policyconfig -> edit -> antivirus` command to specify that the original message is to be included as an RFC 822 attachment for Repaired, Unscannable, Encrypted, or Virus Positive messages. See Send custom alert notification (to recipient only), page 8-12 for more information.

Bounce and Encryption Failure Notification Templates

Bounce and encryption failure notification templates are used in basically the same way as notification templates except that they are used with bounce notifications and message encryption failure notifications. You can specify a custom bounce notification to send while editing a bounce profile and a custom message encryption failure notification while editing an encryption profile.

Figure 14-16 shows a bounce notification template specified in a bounce profile.

Figure 14-16 Bounce Notification Example in a Bounce Profile

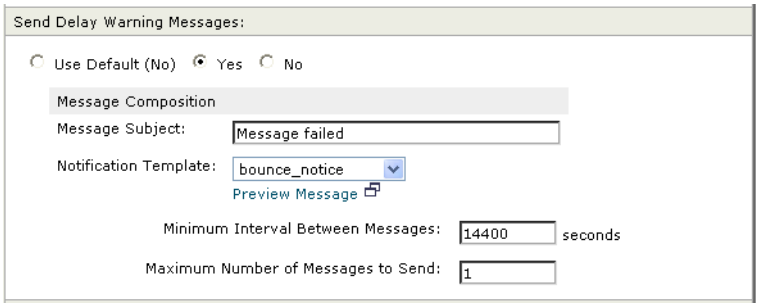

Note You must use RFC-1891 DSNs to use custom templates.

Figure 14-17 shows an encryption failure template specified in an encryption profile.

Figure 14-17 Encryption Failure Notification Example in an Encryption Profile

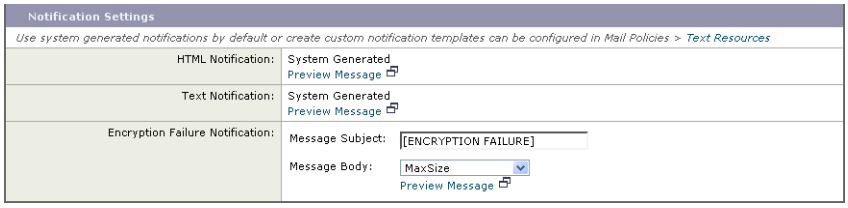

Bounce and Encryption Failure Notification Variables

When creating a bounce or encryption failure notification, you can use any of the notification variables listed in Table 14-5:

Table 14-5 Bounce Notification Variables

Variable	Substituted With
$Subject	The subject of the original message.
$Date	Replaced by the current date, using the format MM/DD/YYYY.
$Time	Replaced by the current time, in the local time zone.
$GMTimeStamp	Replaced by the current time and date, as would be found in the Received: line of an email message, using GMT.
$MID	Replaced by the Message ID, or "MID" used internally to identify the message. Not to be confused with the RFC822 "Message-Id" value (use $Header to retrieve that).
$BouncedRecipient	Bounced recipient address
$BounceReason	Reason for this notification
$remotehost	Replaced by the hostname of the system that sent the message to the Cisco IronPort appliance.

DLP Notification Templates

DLP notification templates are used when you configure your appliance to use the RSA Email DLP feature. The notification informs recipients that an outgoing message may contain sensitive data that violates your organization's data loss prevention policies. You can specify a custom DLP notification while editing a DLP policy in the DLP Policy Manager.

Figure 14-18 shows an example of a DLP notification template being used in a DLP policy.

Figure 14-18 DLP Notification Templates Enabled in a DLP Policy

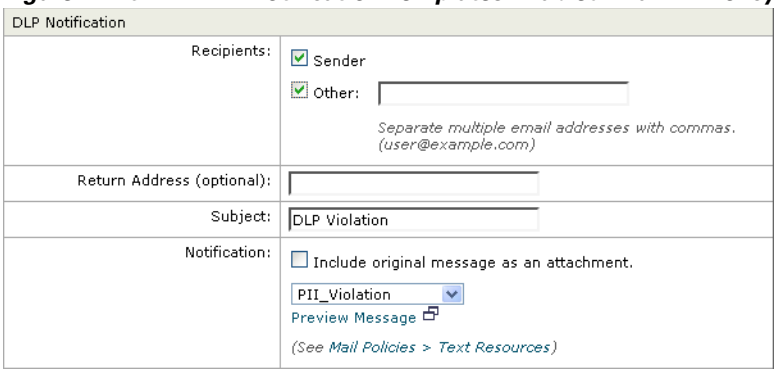

DLP Notification Variables

DLP notification templates can use the following variables.

Table 14-6 DLP Notification Variables

Variable	Substituted With
$DLPPolicy	Replaced by the name of the email DLP policy violated.
$DLPSeverity	Replaced by the severity of violation. Can be “Low,” “Medium,” “High,” or “Critical.”
$DLPRiskFactor	Replaced by the risk factor of the message’s sensitive material (score 0 - 100).
$To	Replaced by the message To: header (not the Envelope Recipient).
$From	Replaced by the message From: header (not the Envelope Sender).
$Subject	Replaced by the subject of the original message.
$Date	Replaced by the current date, using the format MM/DD/YYYY.
$Time	Replaced by the current time, in the local time zone.
$GMTimestamp	Replaced by the current time and date, as would be found in the Received: line of an email message, using GMT.
$MID	Replaced by the Message ID, or “MID” used internally to identify the message. Not to be confused with the RFC822 “Message-Id” value (use $Header to retrieve that).
$Group	Replaced by the name of the sender group the sender matched on when injecting the message. If the sender group had no name, the string “>Unknown<” is inserted.
$Reputation	Replaced by the SenderBase Reputation score of the sender. If there is no reputation score, it is replaced with “None”.
$filenames	Replaced with a comma-separated list of the message’s attachments’ filenames.
$filetypes	Replaced with a comma-separated list of the message's attachments' file types.
$filesizes	Replaced with a comma-separated list of the message’s attachment’s file sizes.

Table 14-6 DLP Notification Variables

Variable	Substituted With
$remotehost	Replaced by the hostname of the system that sent the message to the Cisco IronPort appliance.
$AllHeaders	Replaced by the message headers.
$EnvelopeFrom	Replaced by the Envelope Sender (Envelope From, <MAIL FROM>) of the message.
$Hostname	Replaced by the hostname of the Cisco IronPort appliance.
$bodysize	Replaced by the size, in bytes, of the message.
$header['*string*']	Replaced by the value of the quoted header, if the original message contains a matching header. Note that double quotes may also be used.
$remoteip	Replaced by the IP address of the system that sent the message to the Cisco IronPort appliance.
$recvlistener	Replaced by the nickname of the listener that received the message.
$dropped_filenames	Same as `$filenames`, but displays list of dropped files.
$dropped_filename	Returns only the most recently dropped filename.
$recvint	Replaced by the nickname of the interface that received the message.
$timestamp	Replaced by the current time and date, as would be found in the Received: line of an email message, in the local time zone.
$Time	Replaced by the current time, in the local time zone.
$orgid	Replaced by the SenderBase Organization ID (an integer value).
$enveloperecipients	Replaced by all Envelope Recipients (Envelope To, <RCPT TO>) of the message.
$dropped_filetypes	Same as `$filetypes`, but displays list of dropped file types.
$dropped_filetype	Returns only the file type of the most recently dropped file.

Encryption Notification Templates

Encryption notification templates are used when you configure Cisco IronPort Email Encryption to encrypt outbound email. The notification informs recipients that they have received an encrypted message and provides instructions for reading it. You can specify a custom encryption notification to send with encrypted messages. You specify both an HTML and a text encryption notification when you create an encryption profile. Therefore, if you want to create a custom profile, you should create both text and HTML notifications.

Figure 14-19 shows encryption notifications specified in an encryption profile.

Figure 14-19 Encryption Notification Templates Enabled on Encryption Profile

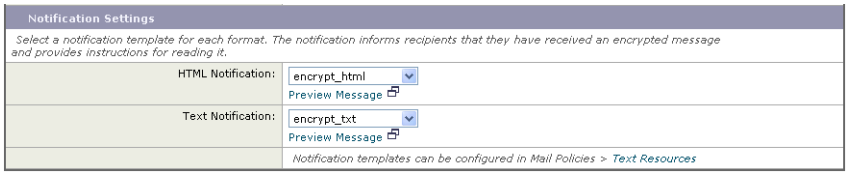

CHAPTER 15

System Administration

System administration in general is handled primarily via the System Administration menu in the Graphical User Interface (GUI). Some system administration features are accessible only via the Command Line Interface (CLI).

In addition, you may want to access the Cisco IronPort appliance's system monitoring features via the Cisco IronPort Graphical User Interface (GUI), which is described in Other Tasks in the GUI, page -441.

Note Several of the features or commands described in this section will affect, or be affected by routing precedence. Please see IP Addresses, Interfaces, and Routing, page B-3 for more information.

Upgrading AsyncOS

Before You Upgrade

Upgrading AsyncOS uses the following two step process:

Step 1 **Configure the upgrade settings.** You can configure settings that affect how the Email Security appliance downloads the upgrade information. For example, you can choose where to download the upgrade images from and more. For more information, see Configuring Upgrade Settings from the GUI, page 15-6.

Step 2 **Upgrade AsyncOS**. After you configure the upgrade settings, upgrade the version of AsyncOS on the appliance. For more information see Upgrading AsyncOS After Configuring Update Setings, page 15-2 and Upgrading AsyncOS from the CLI, page 15-3.

As a best practice, Cisco recommends preparing for an upgrade by taking the following steps:

Step 1 Save the XML config file off-box.

Step 2 If you are using the Safelist/Blocklist feature, export the list off-box.

Step 3 Suspend all listeners. If you perform the upgrade from the CLI, use the `suspendlistener` command. If you perform the upgrade from the GUI, listener suspension occurs automatically.

Step 4 Wait for the queue to empty. You can use the `workqueue` command to view the number of messages in the work queue or the `rate` command in the CLI to monitor the message throughput on your appliance.

Note Re-enable the listeners post-upgrade.

Upgrading AsyncOS After Configuring Update Setings

Step 1 Click **Available Upgrades** on the System Administration > System Upgrade page. The Available Upgrades page is displayed.

Figure 15-1 The Available Upgrades Page

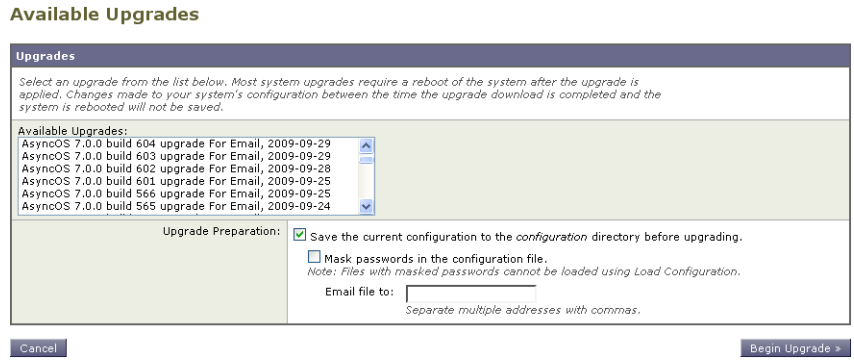

Step 2 Select an upgrade from the list of available upgrades.

Step 3 Choose whether or not to save the current configuration to the `configuration` directory.

Step 4 Choose whether or not to mask the passwords in the configuration file.

Note You cannot load a configuration file with masked passwords using the Configuration File page in the GUI or the `loadconfig` command in the CLI.

Step 5 If you want to email a copies of the configuration file, enter the email addresses to which you want to email the file. Use commas to separate multiple email addresses.

Step 6 Click **Begin Upgrade**. A progress bar appears near the top of the page. You may be asked one or more times to confirm changes or read and agree to new license agreements, etc.

Step 7 After the upgrade is finished, you are asked to reboot the appliance.

Step 8 Click **Reboot Now**.

Upgrading AsyncOS from the CLI

Issue the `upgrade` command to show a list of available upgrades. Select the desired upgrade from the list to install it. You may be asked to confirm messages or read and agree to license agreements, etc. You can choose whether or not to save the current configuration to the `configuration` directory. If so, you can choose whether or not to mask the passwords in the configuration file. You can also choose to email a copy of the configuration files.

Note You cannot load a configuration file with masked passwords using the `loadconfig` command.

When upgrading, do not pause for long amounts of time at the various prompts. If the TCP session should happen to time out during the download, your upgrade may fail.

Differences from Traditional Upgrading Method

Please note these differences when upgrading AsyncOS from a local server as opposed to the traditional method:

Step 1 The upgrading installs immediately *while downloading*.

A banner displays for 10 seconds at the beginning of the upgrade process. While this banner is displayed, you have the option to type Control-C to exit the upgrade process before downloading starts.

Configuring AsyncOS Upgrade Settings

You can configure how the Email Security appliance downloads AsyncOS upgrades. Cisco IronPort provides two methods (or "sources") for upgrading: streaming upgrades and remote upgrades.

With streaming upgrades, your Cisco IronPort appliances download the AsyncOS upgrades directly from the Cisco IronPort update servers. Each Cisco IronPort appliance downloads the upgrade separately. For more information, see Streaming Upgrade Overview, page 15-4.

For remote upgrades, your Cisco IronPort appliances download the AsyncOS upgrades from a server within your network. You only download the upgrade image from Cisco IronPort one time, and then serve it to your Cisco IronPort appliances. For more information, see Remote Upgrade Overview, page 15-5.

Use the Security Services > Service Updates page to switch between the two upgrading methods (streaming is the default), as well as configure the interface to use to download the upgrade and proxy server settings. For more information, see Configuring Upgrade Settings from the GUI, page 15-6. Optionally, use the `updateconfig` command in the CLI.

Figure 15-2 The Service Updates Page

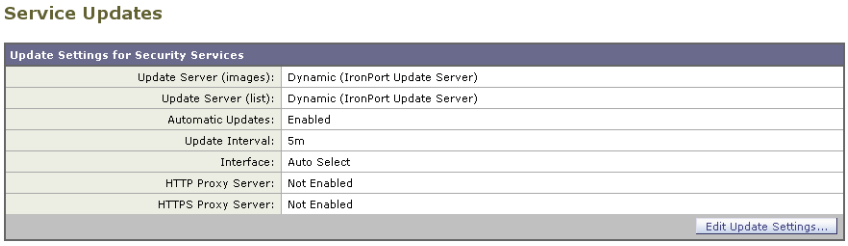

Note Regardless of which upgrade method you use, you should also consider saving your configuration via the `saveconfig` command after your upgrade is complete. For more information, see "Managing the Configuration File" in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.

Upgrading Clustered Systems

If you are upgrading clustered machines, please see "Upgrading Machines in a Cluster" in the "Centralized Management" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Streaming Upgrade Overview

In Streaming upgrades, your Cisco IronPort appliance connects directly to the Cisco IronPort update servers to find and download upgrades:

Figure 15-3 Streaming Update Method

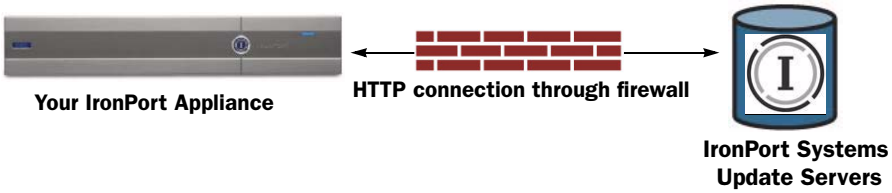

Cisco IronPort Systems uses a distributed upgrade server architecture to make sure customers can quickly download AsyncOS upgrades wherever in the world they are located. Because of this distributed server architecture, the Cisco IronPort update servers use dynamic IP addresses. If you have strict firewall policies, you may need to configure a static location for AsyncOS upgrades. For more information, see Configuring a Static Address for Streaming Upgrades, page 15-4.

You will need to create a firewall rule to allow downloading of upgrades from Cisco IronPort update servers on ports 80 and 443. If you have any existing firewall rules allowing download of legacy upgrades from `upgrades.ironport.com` on ports such as 22, 25, 80, 4766, they will need to be removed and/or replaced with revised firewall rules. For more information, see Appendix C, "Firewall Information".

Configuring a Static Address for Streaming Upgrades

The McAfee Anti-Virus and Cisco IronPort AsyncOS update servers use dynamic IP addresses. If you have strict firewall policies, you may need to configure a static location for updates and AsyncOS upgrades.

Step 1 Contact Cisco IronPort Customer support to obtain the static URL address.

Step 2 Create a firewall rule to allow downloading of upgrades from the static IP address on port 80.

Step 3 Navigate to the Security Services > Service Updates page, and click Edit **Update Settings**.

Step 4 On the Edit Update Settings page, in the "Update Servers (images)" section, choose Local Update Servers and enter the static URL received in step 1 in the Base URL field for AsyncOS upgrades and McAfee Anti-Virus definitions.

Step 5 Verify that IronPort Update Servers is selected for the "Update Servers (list)" section.

Step 6 Submit and commit your changes.

Remote Upgrade Overview

You can also download AsyncOS upgrade images to a local server and host upgrades from within your own network rather than obtaining upgrades directly from Cisco IronPort's update servers. Using this feature, an upgrade image is downloaded via HTTP to any server in your network that has access to the Internet. If you choose to download the update image, you can then configure an internal HTTP server (an "update manager") to host the AsyncOS images to your Cisco IronPort appliances.

Figure 15-4 Remote Update Method

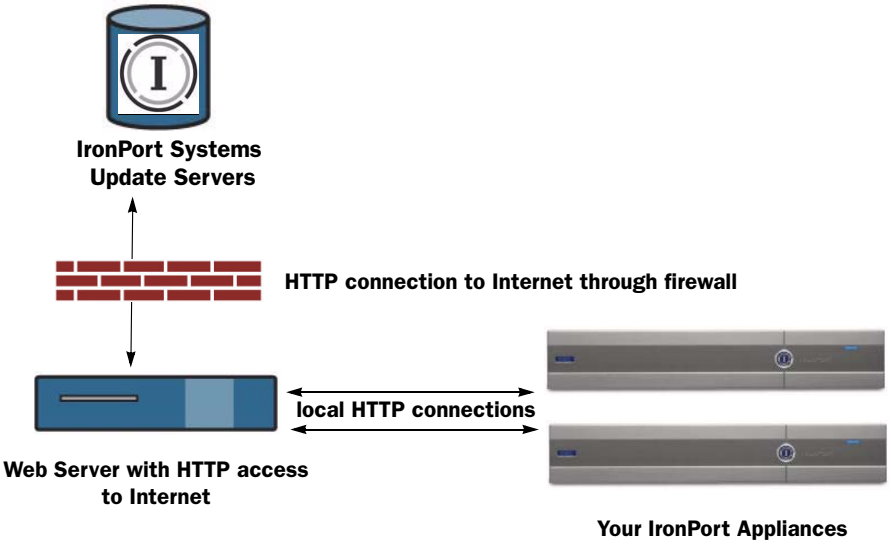

Step 1 Configure a local server to retrieve and serve the upgrade files.

Step 2 Download the upgrade files.

Step 3 Configure the appliance to use the local server using either the Security Services > Service Updates page in the GUI or the `updateconfig` command in the CLI.

Step 4 Upgrade the appliance using either the System Administration > System Upgrade page or the `upgrade` command in the CLI.

Hardware and Software Requirements for Remote Upgrades

For *downloading* AsyncOS upgrade files, you must have a system in your internal network that has:

- Internet access to the Cisco IronPort Systems update servers.
- A web browser (see Browser Requirements, page 2-2).

Note For this release, if you need to configure a firewall setting to allow HTTP access to this address, you must configure it using the DNS name and not a specific IP address.

For *hosting* AsyncOS update files, you must have a server in your internal network that has:

- A web server — for example, Microsoft IIS (Internet Information Services) or the Apache open source server — which:
  - supports the display of directory or filenames in excess of 24 characters
  - has directory browsing enabled
  - is configured for anonymous (no authentication) or basic ("simple") authentication
  - contains at least 350MB of free disk space for each AsyncOS update image

Hosting a Remote Upgrade Image

After setting up a local server, go to `http://updates.ironport.com/fetch_manifest.html` to download a ZIP file of an upgrade image. To download the image, enter your serial number and the version number of the Cisco IronPort appliance. You will then be presented with a list of available upgrades. Click on the upgrade version that you want to download, and unzip the ZIP file in the root directory on the local server while keeping the directory structure intact. To use the upgrade image, configure the appliance to use the local server on the Edit Update Settings page (or use `updateconfig` in the CLI).

The local server also hosts an XML file that limits the available AsyncOS upgrades for the Cisco IronPort appliances on your network to the downloaded upgrade image. This file is called the "manifest." The manifest is located in the `asyncos` directory of the upgrade image ZIP file. After unzipping the ZIP file in the root directory of the local server, enter the full URL for the XML file, including the filename, on the Edit Update Settings page (or use `updateconfig` in the CLI).

For more information about remote upgrades, please see the Cisco IronPort Knowledge Base or contact your Cisco IronPort Support provider.

Note Only use a local update server for AsyncOS upgrade images, not security update images. When you specify a local update server, the local server does not automatically receive updated security updates from Cisco IronPort, so the appliances in your network eventually become out of date. Use a local update server for upgrading AsyncOS, and then change the update and upgrade settings back to use the Cisco IronPort update servers so the security services update automatically again.

Configuring Upgrade Settings from the GUI

Update settings include the source for the AsyncOS upgrade (remote or streaming), the interface to use to download the upgrade, and proxy server settings. In addition to AsyncOS upgrades, you can also edit settings for various Cisco IronPort services such as anti-spam, ant-virus, and Outbreak Filter services. For information about updating services, see Service Updates, page 15-10.

Step 1 Click Edit Update Settings on the Security Services > Service Updates page.

The Edit Update Settings page is displayed.

Step 2 Choose whether to download the AsyncOS upgrade image from the Cisco IronPort update servers or a local server.

If you choose a local server, enter the base URL for the local server hosting the AsyncOS upgrade image. If the server requires authentication, you can also enter a valid user name and password.

Note When you specify a local server for AsyncOS upgrades, the local server does not automatically receive updated McAfee Anti-Virus definitions, so the appliances in your network eventually become out of date. Change the settings back to use the Cisco IronPort update servers after the upgrade so the McAfee Anti-Virus definitions update automatically again.

Step 3 If you choose to download the AsyncOS upgrade image from a local server, select the local server as the source for the list of available updates (the manifest XML file). Enter the full URL for the manifest, including the file name, and the HTTP port number. If the server requires authentication, you can also enter a valid user name and password.

Step 4 Select the interface to use for the upgrade.

Step 5 Enter HTTP proxy server or HTTPS proxy server information if desired.

Step 6 Submit and commit changes.

Configuring Upgrade Settings from the CLI

To tell your appliances where to retrieve the AsyncOS upgrade (local or from Cisco IronPort's servers), run the `updateconfig` command. To install an upgrade, run the `upgrade` command.

The updateconfig Command

The `updateconfig` command is used to tell your Cisco IronPort appliance where to look for service updates, including AsyncOS upgrades. By default, when you type the `upgrade` command, the appliance will contact Cisco IronPort's upgrade servers for the latest update. For remote upgrades, issue the `updateconfig` command and configure the appliance to use a local update server (the local server configured above).

Note You can use the `ping` command to ensure that the appliance can contact the local server. You can also use the `telnet` command to telnet to port 80 of the local server to ensure the local server is listening on that port.

AsyncOS Reversion

AsyncOS includes the ability to revert the AsyncOS operating system to a previous qualified build for emergency uses.

Note After upgrading to AsyncOS 7.0, you cannot revert to a version of AsyncOS earlier than 6.5.

Available Versions

Because upgrades cause one-way transformation of key subsystems, the reversion process is complex and requires qualification by Cisco IronPort Quality Assurance teams. Cisco IronPort certifies specific versions of CASE, Sophos, Outbreak Filters, and McAfee to AsyncOS versions. Not all prior versions of the AsyncOS operating system are available for reversion. The earliest AsyncOS version supported for this functionality is AsyncOS 5.5.0; prior versions of AsyncOS are not supported.

Important Note About Reversion Impact

Using the `revert` command on a Cisco IronPort appliance is a very destructive action. This command destroys all configuration logs and databases. Only the network information for the management interface is preserved--all other network configuration is deleted. In addition, reversion disrupts mail handling until the appliance is reconfigured. Because this command destroys network configuration, you may need physical local access to the Cisco IronPort appliance when you want to issue the `revert` command.

Warning **You must have a configuration file for the version you wish to revert to. Configuration files are *not* backwards-compatible.**

Performing AsyncOS Reversion

Step 1 Ensure that you have the configuration file for the version you wish to revert to. Configuration files are not backwards-compatible. To do this, you can email the file to yourself or FTP the file. A simple way to do this is to run the `mailconfig` CLI command.

Step 2 Save a backup copy of the current configuration of your appliance (with passwords unmasked) on another machine.

Note This is not the configuration file you will load after reverting.

Step 3 If you use the Safelist/Blocklist feature, export the Safelist/Blocklist database to another machine.

Step 4 Wait for the mail queue to empty.

Step 5 Log into the CLI of the appliance you want to revert.

When you run the `revert` command, several warning prompts are issued. After these warning prompts are accepted, the revert action takes place immediately. Therefore, do not begin the reversion process until after you have completed the pre-reversion steps.

Step 6 From the CLI, Issue the `revert` command.

Note The reversion process is time-consuming. It may take fifteen to twenty minutes before reversion is complete and console access to the Cisco IronPort appliance is available again.

The following example shows the `revert` command:

```
mail.mydomain.com> revert


This command will revert the appliance to a previous version of AsyncOS.


WARNING: Reverting the appliance is extremely destructive.

The following data will be destroyed in the process:

- all configuration settings (including listeners)

- all log files

- all databases (including messages in Virus Outbreak and Policy

quarantines)

- all reporting data (including saved scheduled reports)

- all message tracking data

- all IronPort Spam Quarantine message and end-user safelist/blocklist data


Only the network settings will be preserved.


Before running this command, be sure you have:

- saved the configuration file of this appliance (with passwords

unmasked)

- exported the IronPort Spam Quarantine safelist/blocklist database

   to another machine (if applicable)

- waited for the mail queue to empty


Reverting the device causes an immediate reboot to take place.

After rebooting, the appliance reinitializes itself and reboots again to the desired
version.


Do you want to continue?
```

```
Are you *really* sure you want to continue? yes


 Available version                Install date

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

 Available version                Install date

1. 5.5.0-236                        Tue Aug 28 11:03:44 PDT 2007

2. 5.5.0-330                        Tue Aug 28 13:06:05 PDT 2007

3. 5.5.0-418                        Wed Sep  5 11:17:08 PDT 2007


Please select an AsyncOS version: 2

You have selected "5.5.0-330".

The system will now reboot to perform the revert operation.
```

Step 7 The appliance will reboot twice.

Step 8 After the machine reboots twice, use the serial console to configure an interface with an accessible IP address using the `interfaceconfig` command.

Step 9 Enable FTP or HTTP on one of the configured interfaces.

Step 10 Either FTP the XML configuration file you created, or paste it into the GUI interface.

Step 11 Load the XML configuration file of the version you are reverting to.

Step 12 If you use the Safelist/Blocklist feature, import and restore the Safelist/Blocklist database.

Step 13 Commit your changes.

The reverted Cisco IronPort appliance should now run using the selected AsyncOS version.

Service Updates

Many of the settings used to configure how the Cisco IronPort appliance updates various services (such as the anti-spam, anti-virus, and Outbreak Filter services) are accessible via the Service Updates page from the Security Services menu or via the `updateconfig` command in the CLI.

The Service Updates Page

The Service Updates page (available via the Security Services menu in the GUI) displays the current settings for updating various services for your Cisco IronPort appliance. The update settings include: Update Server (images), Update Server (list), Update URLs for various components, Enable Automatic Updates, Automatic Update interval, and the HTTP and HTTPS Proxy Servers.

Note The Cisco IronPort update servers use dynamic IP addresses. If you have strict firewall policies, you may need to configure a static location for security component updates and AsyncOS upgrades. If you determine that your firewall settings require a static IP address for updates and upgrades, follow instructions below for editing the update settings and contact Cisco IronPort Customer support to obtain the required URL addresses.

Editing Update Settings

To edit the update settings for your Cisco IronPort appliance, click the **Edit Update Settings** button. You can configure the following types of settings: Update Servers (images), Update Servers (list), Automatic Updates, Interface, and Proxy Servers. See Table 15-1 on page 15-13 for more details on the update settings.

Figure 15-5 shows the settings available for Update Servers.

Figure 15-5 Update Servers Settings for Images and Lists

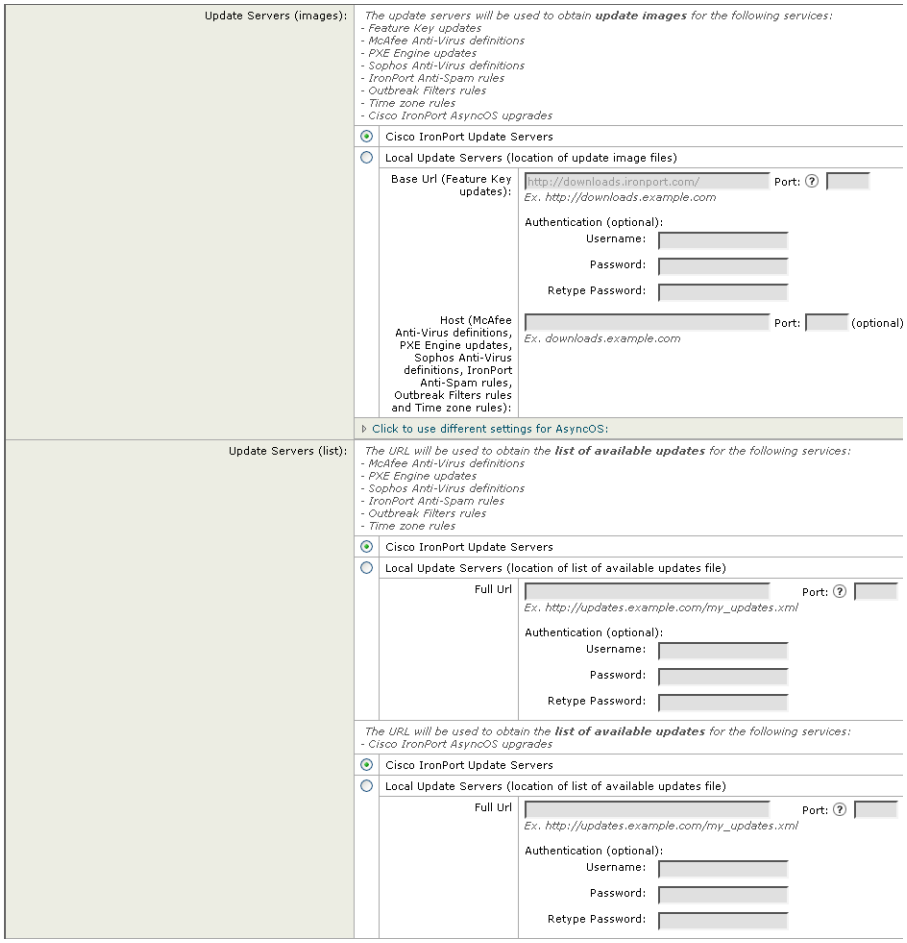

Figure 15-6 shows the settings available for Automatic Updates and the Interface.

Figure 15-6 Automatic Updates and Interfaces Settings

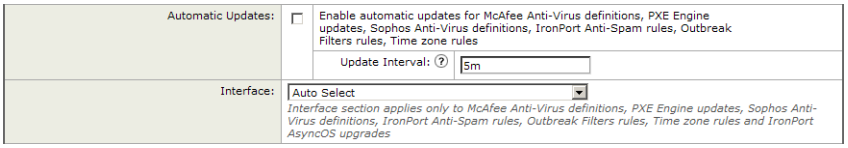

Figure 15-7 shows the settings available for Proxy Servers.

Figure 15-7 Proxy Servers Settings

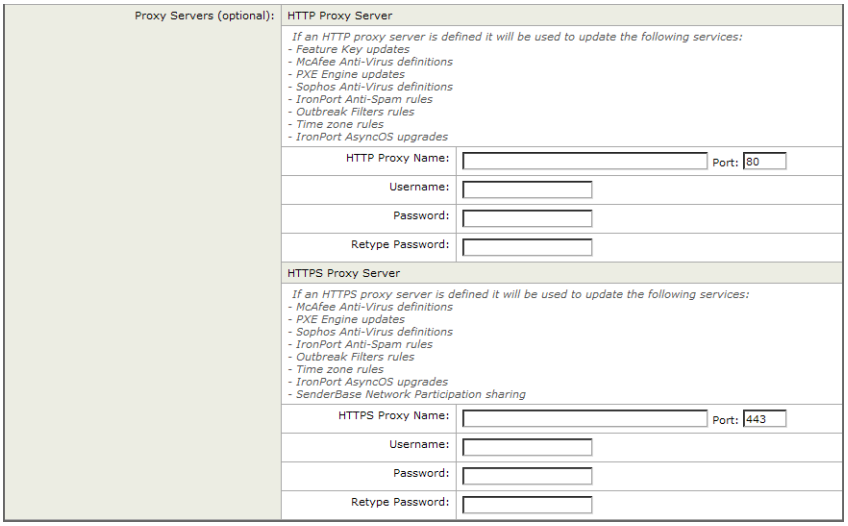

Table 15-1 describes the update settings you can configure.

Table 15-1 Update Settings

Setting	Description
Update Servers (images)	Choose whether to download Cisco IronPort AsyncOS upgrade images and service updates from the Cisco IronPort update servers or a from a local web server. The default is the Cisco IronPort update servers for both upgrades and updates. In addition to AsyncOS upgrades, these servers are used to obtain update images for Sophos and McAfee Anti-Virus definitions, Cisco IronPort Anti-Spam and Cisco IronPort Intelligent Multi-Scan rules, Outbreak Filter rules, time zone rules, feature key updates, and PXE Engine updates. To view settings for AsyncOS upgrades, click the **Click to use different settings for AsyncOS** link. You might want to choose a local web server to: • Download images from Cisco IronPort, if you need to enter a static address provided by Cisco IronPort Customer Support. • Download Cisco IronPort AsyncOS upgrade images at your convenience. (You can still download service update images dynamically from the Cisco IronPort update servers.) When you choose a local update server, enter the base URL and port number for the servers used to download the upgrades and updates. If the server requires authentication, you can also enter a valid user name and password. **Note** Cisco IronPort Intelligent Multi-Scan requires a second local server to download updates for third-party anti-spam rules.
Update Servers (lists)	Choose whether to download the lists of available upgrades and service updates (the manifest XML files) from the Cisco IronPort update servers or from a local web server. The manifest XML file includes updates for different security components, such as McAfee Anti-Virus and the PXE Engine, as well as AsyncOS upgrades. The default for both upgrades and updates is the Cisco IronPort update servers. There are separate sections for specifying servers for updates and for AsyncOS upgrades. If you choose local update servers, enter the full path to the manifest XML file for each list including the file name and port number for the server. If you leave the port field blank, AsyncOS uses port 80. If the server requires authentication, you can also enter a valid user name and password. For more information, see Remote Upgrade Overview, page 15-5.
Automatic Updates	Enable automatic updates and the update interval (how often the appliance checks for updates) for Sophos and McAfee Anti-Virus definitions, Cisco IronPort Anti-Spam rules, Cisco IronPort Intelligent Multi-Scan rules, PXE Engine updates, Outbreak Filter rules, and time zone rules.
Interface	Choose which network interface to use when contacting the update servers for the listed security component updates and Cisco IronPort AsyncOS upgrades. The available proxy data interfaces are shown. By default, the appliance selects an interface to use.

Table 15-1 Update Settings (Continued)

Setting	Description
HTTP Proxy Server	An optional proxy server used for the services listed in the GUI. Note that if you specify a proxy server, it will be used for ALL of these services.
HTTPS Proxy Server	An optional proxy server using HTTPS. If you define the HTTPS proxy server, it will be used to update the services listed in the GUI.

Configuring the Update Server

Step 1 Select either the Cisco IronPort update servers or local update servers for obtaining update images for services

Note If you select a local server as an upgrade source, automatic updates for several security component updates, such as Sophos and McAfee Anti-Virus definitions, cease. To continue updating these security component updates, host the update images or a list of the updates on the local server.

Step 2 If you select local update servers for update images, first enter the base URL, port number, and the optional authentication information for the local server hosting all service updates except AsyncOS upgrades and McAfee Anti-Virus definitions.

Step 3 If you want to get AsyncOS upgrade and McAfee Anti-Virus definitions from a different location than the other service updates, click the Click to use different settings for AsyncOS link at the bottom of the same section and select one of the following options:

- **Cisco IronPort Update Servers**.
- **Local Update Servers**. Enter the base URL and port number for the local server hosting the AsyncOS upgrades.

Step 4 If you select a local update server for the list of available upgrades, enter the full path to the XML file for the list, including the file name, and the HTTP port number as well as the optional authentication information.

Configuring Automatic Updates

Step 1 Select the check box to enable automatic updates.

Step 2 Enter an update interval (time to wait between checks for updates). Add a trailing `m` for minutes and `h` for hours. The maximum update interval is 1 hour.

Specify an HTTP Proxy Server (Optional)

Step 1 Enter a server URL and port number.

Step 2 Enter a username and password for an account on that server, if necessary.

Step 3 Submit and commit your changes.

Specify an HTTPS Proxy Server (Optional)

Step 1 Enter a server URL and port number.

Step 2 Enter a username and password for an account on that server, if necessary.

Step 3 Submit and commit your changes.

Configuring the Return Address for Various Generated Messages

You can configure the envelope sender for mail generated by AsyncOS for the following circumstances:

- Anti-Virus notifications
- Bounces
- Notifications (`notify()` and `notify-copy()` filter actions)
- Quarantine notifications (and "Send Copy" in quarantine management)
- Reports

You can specify the display, user, and domain names of the return address. You can also choose to use the Virtual Gateway domain for the domain name.

Use the Return Addresses page available on the System Administration menu in the GUI, or use the `addressconfig` command via the CLI.

Figure 15-8 The Return Addresses Page

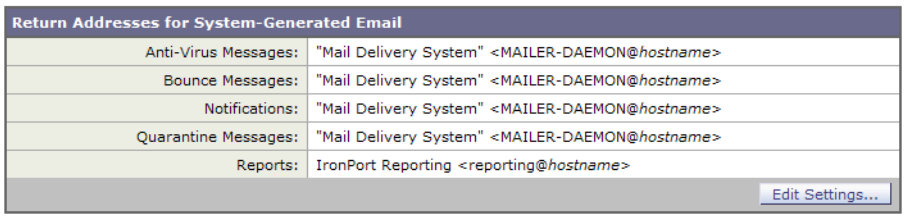

To modify the return address for system-generated email messages via the GUI, click **Edit Settings** on the Return Addresses page. Make changes to the address or addresses you want to modify, click **Submit**, and, finally, commit your changes.

Alerts

Alerts are email notifications containing information about events occurring on the Cisco IronPort appliance. These events can be of varying levels of importance (or severity) from minor to major and pertain generally to a specific component or feature on your appliance. Alerts are generated by the Cisco IronPort appliance. You can specify, at a much more granular level, which alert messages are sent to which users and for which severity of event they are sent. Manage alerts via the System Administration > Alerts page in the GUI (or via the `alertconfig` command in the CLI).

Alerting Overview

The alerting feature consists of two main parts:

- **Alerts** - consist of an Alert **Recipient** (email addresses for receiving alerts), and the alert notification (severity and alert type) sent to the recipient.
- **Alert Settings** - specify global behavior for the alerting feature, including alert sender (FROM:) address, seconds to wait between sending duplicate alerts, and whether to enable AutoSupport (and optionally send weekly AutoSupport reports).

Alerts: Alert Recipients, Alert Classifications, and Severities

Alerts are email messages or notifications containing information about a specific function (or alert classification) or functions such as a hardware or anti-virus problem, sent to an alert recipient. An alert recipient is simply an email address to which the alert notifications are sent. The information contained in the notification is determined by an alert classification and a severity. You can specify which alert classifications, at which severity, are sent to any alert recipient. The alerting engine allows for granular control over which alerts are sent to which alert recipients. For example, you can configure the system to send only specific alerts to an alert recipient, configuring an alert recipient to receive notifications only when Critical (severity) information about the System (alert type) is sent. You can also configure general settings (see Configuring Alert Settings, page 15-21).

See Alert Listing, page 15-22 for a complete list of alerts.

Alert Classifications

AsyncOS sends the following alert classifications:

- System
- Hardware
- Updater
- Outbreak Filters
- Anti-Virus
- Anti-Spam
- Directory Harvest Attack Prevention

Severities

Alerts can be sent for the following severities:

- Critical: Requires immediate attention.
- Warning: Problem or error requiring further monitoring and potentially immediate attention.
- Information: Information generated in the routine functioning of this device.

Alert Settings

Alert settings control the general behavior and configuration of alerts, including:

- The RFC 2822 Header From: when sending alerts (enter an address or use the default “alert@<hostname>”). You can also set this via the CLI, using the `alertconfig -> from` command.

- The initial number of seconds to wait before sending a duplicate alert.
- The maximum number of seconds to wait before sending a duplicate alert.
- The status of AutoSupport (enabled or disabled).
- The sending of AutoSupport's weekly status reports to alert recipients set to receive System alerts at the Information level.

Sending Duplicate Alerts

You can specify the initial number of seconds to wait before AsyncOS will send a duplicate alert. If you set this value to 0, duplicate alert summaries are not sent and instead, all duplicate alerts are sent without any delay (this can lead to a large amount of email over a short amount of time). The number of seconds to wait between sending duplicate alerts (alert interval) is increased after each alert is sent. The increase is the number of seconds to wait plus twice the last interval. So a 5 second wait would have alerts sent at 5 seconds, 15, seconds, 35 seconds, 75 seconds, 155 seconds, 315 seconds, etc.

Eventually, the interval could become quite large. You can set a cap on the number of seconds to wait between intervals via the maximum number of seconds to wait before sending a duplicate alert field. For example, if you set the initial value to 5 seconds, and the maximum value to 60 seconds, alerts would be sent at 5 seconds, 15 seconds, 35 seconds, 60 seconds, 120 seconds, etc.

SMTP Routes and Alerts

Alerts sent from the appliance to addresses specified in the Alert Recipient follow SMTP routes defined for those destinations.

Cisco IronPort AutoSupport

To allow Cisco to better support and design future system changes, the Cisco IronPort appliance can be configured to send Cisco Systems a copy of all alert messages generated by the system. This feature, called AutoSupport, is a useful way to allow our team to be proactive in supporting your needs. AutoSupport also sends weekly reports noting the uptime of the system, the output of the `status` command, and the AsyncOS version used.

By default, alert recipients set to receive Information severity level alerts for System alert types will receive a copy of every message sent to Cisco. This can be disabled if you do not want to send the weekly alert messages internally. To enable or disable this feature, see Configuring Alert Settings, page 15-21.

Alert Messages

Alert messages are standard email messages. You can configure the Header From: address, but the rest of the message is generated automatically.

Alert From Address

You can configure the Header From: address via the **Edit Settings** button or via the CLI (see the *Cisco IronPort AsyncOS CLI Reference Guide*).

Alert Subject

An alert email message's subject follows this format:

```
Subject: [severity]-[hostname]: ([class]) short message
```

Alert Delivery

Because alert messages can be used to inform you of problems within your Cisco IronPort appliance, they are not sent using AsyncOS's normal mail delivery system. Instead, alert messages pass through a separate and parallel email system designed to operate even in the face of significant system failure in AsyncOS.

The alert mail system does not share the same configuration as AsyncOS, which means that alert messages may behave slightly differently from other mail delivery:

- Alert messages are delivered using standard DNS MX and A record lookups.
  - They do not use `smtproutes` in AsyncOS versions older then 5.X.
  - They do cache the DNS entries for 30 minutes and the cache is refreshed every 30 minutes, so in case of DNS failure the alerts still go out.
- Alert messages do not pass through the work queue, so they are not scanned for viruses or spam. They are also not subjected to message filters or content filters.
- Alert messages do not pass through the delivery queue, so they are not affected by bounce profiles or destination control limits.

Example Alert Message

```
Date: 23 Mar 2005 21:10:19 +0000

To: joe@example.com

From: IronPort C60 Alert [alert@example.com]

Subject: Critical-example.com: (Anti-Virus) update via http://newproxy.example.com
failed



The Critical message is:



update via http://newproxy.example.com failed



Version: 4.5.0-419

Serial Number: XXXXXXXXXXXX-XXXXXXX

Timestamp: Tue May 10 09:39:24 2005



For more information about this error, please see

    http://support.ironport.com

If you desire further information, please contact your support provider.
```

Managing Alert Recipients

Log in to the Graphical User Interface (GUI) and click the System Administration tab. (For information about how to access the GUI, see Accessing the GUI, page 2-2.) Click the Alerts link in the left menu.

Figure 15-9 The Alerts Page

Alerts

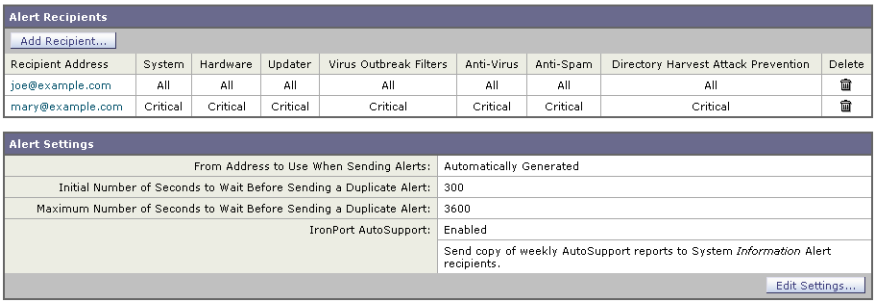

Note If you enabled AutoSupport during System Setup, the email address you specified will receive alerts for all severities and classes by default. You can change this configuration at any time.

The Alerts page lists the existing alert recipients and alert settings.

From the Alerts page, you can:

- Add, configure, or delete alert recipients
- Modify the alert settings

Adding New Alert Recipients

Step 1 Click **Add Recipient** on the Alerts page. The Add Alert Recipients page is displayed:

Figure 15-10 Adding a New Alert Recipient

Add Alert Recipient

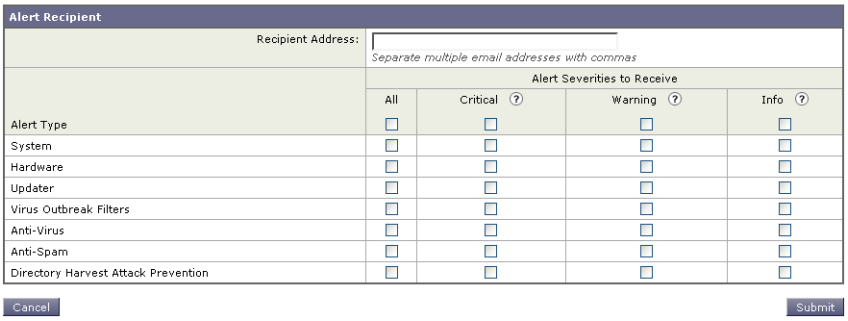

Step 2 Enter the recipient's email address. You can enter multiple addresses, separated by commas.

Step 3 Select which alert severities to receive.

Step 4 Submit and commit your changes.

Configuring Existing Alert Recipients

Step 1 Click the alert recipient in the Alert Recipients listing. The Configure Alert Recipient page is displayed.

Step 2 Make changes to the alert recipient.

Step 3 Submit and commit your changes.

Deleting Alert Recipients

Step 1 Click the trash can icon corresponding to the alert recipient in the Alert Recipient listing.

Step 2 Confirm the deletion by clicking **Delete** in the warning dialog that appears.

Step 3 Commit your changes.

Configuring Alert Settings

Alert settings are global settings, meaning that they affect how all of the alerts behave.

Editing Alert Settings

Step 1 Click **Edit Settings** on the Alerts page. The Edit Alert Settings page is displayed:

Figure 15-11 Editing Alert Settings

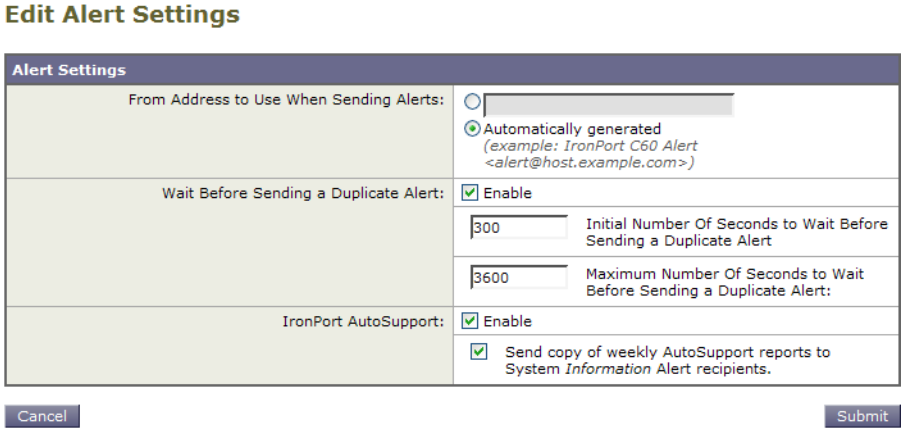

Step 2 Enter a Header From: address to use when sending alerts, or select Automatically Generated ("alert@<hostname>").

Step 3 Mark the checkbox if you want to specify the number of seconds to wait between sending duplicate alerts. For more information, see Sending Duplicate Alerts, page 15-17.

- Specify the initial number of seconds to wait before sending a duplicate alert.
- Specify the maximum number of seconds to wait before sending a duplicate alert.

Step 4 You can enable AutoSupport by checking the IronPort AutoSupport option. For more information about AutoSupport, see Cisco IronPort AutoSupport, page 15-17.

- If AutoSupport is enabled, the weekly AutoSupport report is sent to alert recipients set to receive System alerts at the Information level. You can disable this via the checkbox.

Step 5 Submit and commit your changes.

Alert Listing

The following tables list alerts by classification, including the alert name (internal descriptor used by Cisco IronPort), actual text of the alert, description, severity (critical, information, or warning) and the parameters (if any) included in the text of the message. The value of the parameter is replaced in the actual text of the alert. For example, an alert message below may mention "$ip" in the message text. "$ip" is replaced by the actual IP address when the alert is generated.

Anti-Spam Alerts

Table 15-2 contains a list of the various anti-spam alerts that can be generated by AsyncOS, including a description of the alert and the alert severity.

Table 15-2 Listing of Possible Anti-Spam Alerts

Alert Name	Message and Description	Parameters
AS.SERVER.ALERT	$engine anti-spam - $message $tb Critical. Sent when the anti-spam engine fails.	'**engine**' - The type of anti-spam engine. '**message**' - The log message. '**tb**' - Traceback of the event.
AS.TOOL.INFO_ALERT	Update - $engine - $message Information. Sent when there is a problem with the anti-spam engine.	'**engine**' - The anti-spam engine name '**message**' - The message
AS.TOOL.ALERT	Update - $engine - $message Critical. Sent when an update is aborted due to a problem with one of the tools used to manage the anti-spam engine.	'**engine**' - The anti-spam engine name '**message**' - The message

Anti-Virus Alerts

Table 15-3 contains a list of the various Anti-Virus alerts that can be generated by AsyncOS, including a description of the alert and the alert severity.

Table 15-3 Listing of Possible Anti-Virus Alerts

Alert Name	Message and Description	Parameters
AV.SERVER.ALERT / AV.SERVER.CRITICAL	$engine antivirus - $message $tb Critical. Sent when there is a critical problem with the anti-virus scanning engine.	'**engine**' - The type of anti-virus engine. '**message**' - The log message. '**tb**' - Traceback of the event.

Table 15-3 Listing of Possible Anti-Virus Alerts (Continued)

Alert Name	Message and Description	Parameters
AV.SERVER.ALERT.INFO	$engine antivirus - $message $tb	'**engine**' - The type of anti-virus engine. '**message**' - The log message. '**tb**' - Traceback of the event.
	Information. Sent when an informational event occurs with the anti-virus scanning engine.	
AV.SERVER.ALERT.WARN	$engine antivirus - $message $tb	'**engine**' - The type of anti-virus engine. '**message**' - The log message. '**tb**' - Traceback of the event.
	Warning. Sent when there is a problem with the anti-virus scanning engine.	
MAIL.ANTIVIRUS. ERROR_MESSAGE	MID $mid antivirus $what error $tag	'**mid**' - MID '**what**' - The error that happened. '**tag**' - Virus outbreak name if set.
	Critical. Sent when anti-virus scanning produces an error while scanning a message.	
MAIL.SCANNER. PROTOCOL_MAX_RETRY	MID $mid is malformed and cannot be scanned by $engine.	'**mid**' - MID '**engine**' - The engine being used
	Critical. The scanning engine attempted to scan the message unsuccessfully because the message is malformed. The maximum number of retries has been exceeded, and the message will be processed without being scanned by this engine.	

Directory Harvest Attack Prevention (DHAP) Alerts

Table 15-4 contains a list of the various DHAP alerts that can be generated by AsyncOS, including a description of the alert and the alert severity.

Table 15-4 Listing of Possible Directory Harvest Attack Prevention Alerts

Alert Name	Message and Description	Parameters
LDAP.DHAP_ALERT	LDAP: Potential Directory Harvest Attack detected. See the system mail logs for more information about this attack.	
	Warning. Sent when a possible directory harvest attack is detected.	

Hardware Alerts

Table 15-5 contains a list of the various Hardware alerts that can be generated by AsyncOS, including a description of the alert and the alert severity.

Table 15-5 Listing of Possible Hardware Alerts

Alert Name	Message and Description	Parameters
INTERFACE.ERRORS	Port $port: has detected $in_err input errors, $out_err output errors, $col collisions please check your media settings. Warning. Sent when interface errors are detected.	'**port**' - Interface name. '**in_err**' - The number of input errors since the last message. '**out_err**' - The number of output errors since the last message. '**col**' - The number of packet collisions since the last message.
MAIL.MEASUREMENTS_ FILESYSTEM	The $file_system partition is at $capacity% capacity Warning. Sent when a disk partition is nearing capacity (75%).	'**file_system**' - The name of the filesystem '**capacity**' - How full the filesystem is in percent.
MAIL.MEASUREMENTS_ FILESYSTEM.CRITICAL	The $file_system partition is at $capacity% capacity Critical. Sent when a disk partition reaches 90% capacity (and at 95%, 96%, 97%, etc.).	'**file_system**' - The name of the filesystem '**capacity**' - How full the filesystem is in percent.
SYSTEM.RAID_EVENT_ ALERT	A RAID-event has occurred: $error Warning. Sent when a critical RAID-event occurs.	'**error**' - The text of the RAID error.
SYSTEM.RAID_EVENT_ ALERT_INFO	A RAID-event has occurred: $error Information. Sent when a RAID-event occurs.	'**error**' - The text of the RAID error.

Cisco IronPort Spam Quarantine Alerts

Table 15-6 contains a list of the various Cisco IronPort Spam Quarantine alerts that can be generated by AsyncOS, including a description of the alert and the alert severity.

Table 15-6 Listing of Possible Cisco IronPort Spam Quarantine Alerts

Alert Name	Message and Description	Parameters
ISQ.CANNOT_CONNECT_OFF_ BOX	ISQ: Could not connect to off-box quarantine at $host:$port Information. Sent when AsyncOS was unable to connect to the (off-box) IP address.	'**host**' - address of off-box quarantine '**port**' - port to connect to on off-box quarantine

Table 15-6 Listing of Possible Cisco IronPort Spam Quarantine Alerts (Continued)

Alert Name	Message and Description	Parameters
ISQ.CRITICAL	ISQ: $msg	'**msg**' - message to be displayed
	Critical. Sent when a critical error with Cisco IronPort Spam Quarantine is encountered.	
ISQ.DB_APPROACHING_ FULL	ISQ: Database over $threshold% full	'**threshold**' - the percent full threshold at which alerting begins
	Warning. Sent when the Cisco IronPort Spam Quarantine database is nearly full.	
ISQ.DB_FULL	ISQ: database is full	
	Critical. Sent when the Cisco IronPort Spam Quarantine database is full.	
ISQ.MSG_DEL_FAILED	ISQ: Failed to delete MID $mid for $rcpt: $reason	'**mid**' - MID '**rcpt**' - Recipient or "all" '**reason**' - Why the message was not deleted
	Warning. Sent when an email is not successfully deleted from the Cisco IronPort Spam Quarantine.	
ISQ.MSG_NOTIFICATION_ FAILED	ISQ: Failed to send notification message: $reason	'**reason**' - Why the notification was not sent
	Warning. Sent when a notification message is not successfully sent.	
ISQ.MSG_QUAR_FAILED		
	Warning. Sent when a message is not successfully quarantined.	
ISQ.MSG_RLS_FAILED	ISQ: Failed to release MID $mid to $rcpt: $reason	'**mid**' - MID '**rcpt**' - Recipient or "all" '**reason**' - Why the message was not released
	Warning. Sent when a message is not successfully released.	
ISQ.MSG_RLS_FAILED_ UNK_RCPTS	ISQ: Failed to release MID $mid: $reason	'**mid**' - MID '**reason**' - Why the message was not released
	Warning. Sent when a message is not successfully released because the recipient is unknown.	
ISQ.NO_EU_PROPS	ISQ: Could not retrieve $user's properties. Setting defaults	'**user**' - end user name
	Information. Sent when AsyncOS is unable to retrieve information about a user.	

Table 15-6 Listing of Possible Cisco IronPort Spam Quarantine Alerts (Continued)

Alert Name	Message and Description	Parameters
ISQ.NO_OFF_BOX_HOST_SET	ISQ: Setting up off-box ISQ without setting host	
	Information. Sent when AsyncOS is configured to reference an external quarantine, but the external quarantine is not defined.	

Safelist/Blocklist Alerts

contains a list of the various Safelist/Blocklist alerts that can be generated by AsyncOS, including a description of the alert and the alert severity

Table 15-7 Listing of Possible Safelist/Blocklist Alerts

Alert Name	Message and Description	Parameters
SLBL.DB.RECOVERY_FAILED	SLBL: Failed to recover End-User Safelist/Blocklist database: '$error'.	'**error**' - error reason
	Critical. Failed to recover the Safelist/Blocklist database.	
SLBL.DB.SPACE_LIMIT	SLBL: End-User Safelist/Blocklist database exceeded allowed disk space: $current of $limit.	'**current**' - how much it has used, in MB '**limit**' - the configured limit, in MB
	Critical. The safelist/blocklist database exceeded the allowed disk space.	

System Alerts

Table 15-8 contains a list of the various System alerts that can be generated by AsyncOS, including a description of the alert and the alert severity.

Table 15-8 Listing of Possible System Alerts

Alert Name	Message and Description	Parameters
COMMON.APP_FAILURE	An application fault occurred: $error	'**error**' - The text of the error, typically a traceback.
	Warning. Sent when there is an unknown application failure.	
COMMON.KEY_EXPIRED_ALERT	Your "$feature" key has expired. Please contact your authorized Cisco IronPort sales representative.	'**feature**' - The name of the feature that is about to expire.
	Warning. Sent when a feature key has expired.	

Table 15-8 Listing of Possible System Alerts (Continued)

Alert Name	Message and Description	Parameters
COMMON.KEY_EXPIRING_ALERT	Your "$feature" key will expire in under $days day(s). Please contact your authorized Cisco IronPort sales representative.	'**feature**' - The name of the feature that is about to expire. '**days**' - The number of days it will expire.
	Warning. Sent when a feature key is about to expire.	
COMMON.KEY_FINAL_ EXPIRING_ALERT	This is a final notice. Your "$feature" key will expire in under $days day(s). Please contact your authorized Cisco IronPort sales representative.	'**feature**' - The name of the feature that is about to expire. '**days**' - The number of days it will expire.
	Warning. Sent as a final notice that a feature key is about to expire.	
DNS.BOOTSTRAP_FAILED	Failed to bootstrap the DNS resolver. Unable to contact root servers.	
	Warning. Sent when the appliance is unable to contact the root DNS servers.	
INTERFACE. FAILOVER.FAILURE_ BACKUP_DETECTED	Standby port $port on $pair_name failure	'**port**' - Detected port '**pair_name**' - Failover pair name.
	Warning. Sent when a backup NIC pairing interface fails.	
INTERFACE. FAILOVER.FAILURE_ BACKUP_RECOVERED	Standby port $port on $pair_name okay	'**port**' - Failed port '**pair_name**' - Failover pair name.
	Information. Sent when a NIC pair failover is recovered.	
INTERFACE.FAILOVER. FAILURE_DETECTED	Port $port failure on $pair_name, switching to $port_other	'**port**' - Failed port. '**port_other**' - New port. '**pair_name**' - Failover pair name.
	Critical. Sent when a NIC pairing failover is detected due to an interface failure.	
INTERFACE.FAILOVER. FAILURE_DETECTED_NO_ BACKUP	Port $port_other on $pair_name is down, can't switch to $port_other	'**port**' - Failed port. '**port_other**' - New port. '**pair_name**' - Failover pair name.
	Critical. Sent when a NIC pairing failover is detected due to an interface failure, but a backup interface is not available.	
INTERFACE.FAILOVER. FAILURE_RECOVERED	Recovered network on $pair_name using port $port	'**port**' - Failed port '**pair_name**' - Failover pair name.
	Information. Sent when a NIC pair failover is recovered.	

Table 15-8 Listing of Possible System Alerts (Continued)

Alert Name	Message and Description	Parameters
INTERFACE.FAILOVER. MANUAL	Manual failover to port $port on $pair_name	'**port**' - New active port. '**pair_name**' - Failover pair name.
	Information. Sent when a manual failover to another NIC pair is detected.	
COMMON.INVALID_FILTER	Invalid $class: $error	'**class**' - Either "Filter", "SimpleFilter", etc. '**error**' - Additional why-filter-is-invalid info.
	Warning. Sent when an invalid filter is encountered.	
LDAP.GROUP_QUERY_ FAILED_ALERT	LDAP: Failed group query $name, comparison in filter will evaluate as false	'**name**' - The name of the query.
	Critical. Sent when an LDAP group query fails.	
LDAP.HARD_ERROR	LDAP: work queue processing error in $name reason $why	'**name**' - The name of the query. '**why**' - Why the error happened.
	Critical. Sent when an LDAP query fails completely (after trying all servers).	
LOG.ERROR.*	Critical. Various logging errors.	
MAIL.PERRCPT.LDAP_ GROUP_QUERY_FAILED	LDAP group query failure during per-recipient scanning, possible LDAP misconfiguration or unreachable server.	
	Critical. Sent when an LDAP group query fails during per-recipient scanning.	
MAIL.QUEUE.ERROR.*	Critical. Various mail queue hard errors.	
MAIL.RES_CON_START_ ALERT.MEMORY	This system (hostname: $hostname) has entered a 'resource conservation' mode in order to prevent the rapid depletion of critical system resources. RAM utilization for this system has exceeded the resource conservation threshold of $memory_threshold_start%. The allowed receiving rate for this system will be gradually decreased as RAM utilization approaches $memory_threshold_halt%.	'**hostname**' - The name of the host. '**memory_threshold_start**' - The percent threshold where memory tarpitting starts. '**memory_threshold_halt**' - The percent threshold where the system will halt due to memory being too full.
	Critical. Sent when RAM utilization has exceeded the system resource conservation threshold.	

Table 15-8 Listing of Possible System Alerts (Continued)

Alert Name	Message and Description	Parameters
MAIL.RES_CON_START_ ALERT.QUEUE_SLOW	This system (hostname: $hostname) has entered a 'resource conservation' mode in order to prevent the rapid depletion of critical system resources. The queue is overloaded and is unable to maintain the current throughput.	'**hostname**' - The name of the host.
	Critical. Sent when the mail queue is overloaded and system resource conservation is enabled.	
MAIL.RES_CON_START_ ALERT.QUEUE	This system (hostname: $hostname) has entered a 'resource conservation' mode in order to prevent the rapid depletion of critical system resources. Queue utilization for this system has exceeded the resource conservation threshold of $queue_threshold_start%. The allowed receiving rate for this system will be gradually decreased as queue utilization approaches $queue_threshold_halt%.	'**hostname**' - The name of the host. '**queue_threshold_start**' - The percent threshold where queue tarpitting starts. '**queue_threshold_halt**' - The percent threshold where the system will halt due to the queue being too full.
	Critical. Sent when queue utilization has exceeded the system resource conservation threshold.	
MAIL.RES_CON_START_ ALERT.WORKQ	This system (hostname: $hostname) has entered a 'resource conservation' mode in order to prevent the rapid depletion of critical system resources. Listeners have been suspended because the current work queue size has exceeded the threshold of $suspend_threshold. Listeners will be resumed once the work queue size has dropped to $resume_threshold. These thresholds may be altered via use of the 'tarpit' command on the system CLI.	'**hostname**' - The name of the host. '**suspend_threshold**' - Work queue size above which listeners are suspended. '**resume_threshold**' - Work queue size below which listeners are resumed.
	Information. Sent when listeners are suspended because the work queue size is too big.	
MAIL.RES_CON_START_ ALERT	This system (hostname: $hostname) has entered a 'resource conservation' mode in order to prevent the rapid depletion of critical system resources.	'**hostname**' - The name of the host.
	Critical. Sent when the appliance enters "resource conservation" mode.	

Table 15-8 Listing of Possible System Alerts (Continued)

Alert Name	Message and Description	Parameters
MAIL.RES_CON_STOP_ ALERT	This system (hostname: $hostname) has exited 'resource conservation' mode as resource utilization has dropped below the conservation threshold.	'**hostname**' - The name of the host.
	Information. Sent when the appliance leaves 'resource conservation' mode.	
MAIL.WORK_QUEUE_ PAUSED_NATURAL	work queue paused, $num msgs, $reason	'**num**' - The number of messages in the work queue.
	Critical. Sent when the work queue is paused.	'**reason**' - The reason the work queue is paused.
MAIL.WORK_QUEUE_ UNPAUSED_NATURAL	work queue resumed, $num msgs	'**num**' - The number of messages in the work queue.
	Critical. Sent when the work queue is resumed.	
NTP.NOT_ROOT	Not running as root, unable to adjust system time	
	Warning. Sent when the Cisco IronPort appliance is unable to adjust time because NTP is not running as root.	
QUARANTINE.ADD_DB_ ERROR	Unable to quarantine MID $mid - quarantine system unavailable	'**mid**' - MID
	Critical. Sent when a message cannot be sent to a quarantine.	
QUARANTINE.DB_ UPDATE_FAILED	Unable to update quarantine database (current version: $version; target $target_version)	'**version**' - The schema version detected.
	Critical. Sent when a quarantine database cannot be updated.	'**target_version**' - The target schema version.
QUARANTINE.DISK_ SPACE_LOW	The quarantine system is unavailable due to a lack of space on the $file_system partition.	'**file_system**' - The name of the filesystem.
	Critical. Sent when the disk space for quarantines is full.	
QUARANTINE.THRESHOLD_ALERT	Quarantine "$quarantine" is $full% full	'**quarantine**' - The name of the quarantine.
	Warning. Sent when a quarantine reaches 5%, 50%, or 75% of capacity.	'**full**' - The percentage of how full the quarantine is.
QUARANTINE.THRESHOLD_ALERT .SERIOUS	Quarantine "$quarantine" is $full% full	'**quarantine**' - The name of the quarantine.
	Critical. Sent when a quarantine reaches 95% of capacity.	'**full**' - The percentage of how full the quarantine is.

Table 15-8 Listing of Possible System Alerts (Continued)

Alert Name	Message and Description	Parameters
REPORTD.DATABASE_ OPEN_FAILED_ALERT	The reporting system has encountered a critical error while opening the database. In order to prevent disruption of other services, reporting has been disabled on this machine. Please contact customer support to have reporting enabled. The error message is: $err_msg	'**err_msg**' - The error message raised
	Critical. Sent if the reporting engine is unable to open the database.	
REPORTD.AGGREGATION_DISABL ED_ALERT	Processing of collected reporting data has been disabled due to lack of logging disk space. Disk usage is above $threshold percent. Recording of reporting events will soon become limited and reporting data may be lost if disk space is not freed up (by removing old logs, etc.). Once disk usage drops below $threshold percent, full processing of reporting data will be restarted automatically.	'**threshold**' - The threshold value
	Warning. Sent if the system runs out of disk space. When the disk usage for a log entry exceeds the log usage threshold, reportd disables aggregation and sends the alert.	
REPORTING.CLIENT. UPDATE_FAILED_ALERT	Reporting Client: The reporting system has not responded for an extended period of time ($duration).	'**duration**' - Length of time the client has been trying to contact the reporting daemon. This is a string in a human readable format ('1h 3m 27s').
	Warning. Sent if the reporting engine was unable to save reporting data.	
REPORTING.CLIENT. JOURNAL.FULL	Reporting Client: The reporting system is unable to maintain the rate of data being generated. Any new data generated will be lost.	
	Critical. Sent if the reporting engine is unable to store new data.	
REPORTING.CLIENT. JOURNAL.FREE	Reporting Client: The reporting system is now able to handle new data.	
	Information. Sent when the reporting engine is again able to store new data.	

Table 15-8 Listing of Possible System Alerts (Continued)

Alert Name	Message and Description	Parameters
PERIODIC_REPORTS. REPORT_TASK.BUILD_ FAILURE	A failure occurred while building periodic report '$report_title'. This subscription has been removed from the scheduler.	'**report_title**' - the report title
	Critical. Sent when the reporting engine is unable to build a report.	
PERIODIC_REPORTS. REPORT_TASK.EMAIL_ FAILURE	A failure occurred while emailing periodic report '$report_title'. This subscription has been removed from the scheduler.	'**report_title**' - the report title
	Critical. Sent when a report could not be emailed.	
PERIODIC_REPORTS. REPORT_TASK.ARCHIVE_FAILURE	A failure occurred while archiving periodic report '$report_title'. This subscription has been removed from the scheduler.	'**report_title**' - the report title
	Critical. Sent when a report could not be archived.	
SENDERBASE.ERROR	Error processing response to query $query: response was $response	'**query**' - The query address. '**response**' - Raw data of response received.
	Information. Sent when an error occurred while processing a response from SenderBase.	
SMTPAUTH.FWD_SERVER_FAILED _ALERT	SMTP Auth: could not reach forwarding server $ip with reason: $why	'**ip**' - The IP of the remote server. '**why**' - Why the error happened.
	Warning. Sent when the SMTP Authentication forwarding server is unreachable.	
SMTPAUTH.LDAP_QUERY_FAILED	SMTP Auth: LDAP query failed, see LDAP debug logs for details.	
	Warning. Sent when an LDAP query fails.	
SYSTEM.HERMES_ SHUTDOWN_FAILURE. REBOOT	While preparing to ${what}, failed to stop mail server gracefully: ${error}$what:=reboot	'**error**' - The error that happened.
	Warning. Sent when there was a problem shutting down the system on reboot.	
SYSTEM.HERMES_ SHUTDOWN_FAILURE. SHUTDOWN	While preparing to ${what}, failed to stop mail server gracefully: ${error}$what:=shut down	'**error**' - The error that happened.
	Warning. Sent when there was a problem shutting down the system.	

Table 15-8 Listing of Possible System Alerts (Continued)

Alert Name	Message and Description	Parameters
SYSTEM. RCPTVALIDATION.UPDATE_FAILED	Error updating recipient validation data: $why	'why' - The error message.
	Critical. Sent when a recipient validation update failed.	
SYSTEM.SERVICE_ TUNNEL.DISABLED	Tech support: Service tunnel has been disabled	
	Information. Sent when a tunnel created for Cisco IronPort Support Services is disabled.	
SYSTEM.SERVICE_ TUNNEL.ENABLED	Tech support: Service tunnel has been enabled, port $port	'**port**' - The port used for the service tunnel.
	Information. Sent when a tunnel created for Cisco IronPort Support Services is enabled.	

Updater Alerts

Table 15-9 contains a list of the varius Updater alerts that can be generated by AsyncOS.

Table 15-9 Listing of Possible Updater Alerts

Alert Name	Message and Description	Parameters
UPDATER.APP.UPDATE_ABANDONED	$app abandoning updates until a new version is published. The $app application tried and failed $attempts times to successfully complete an update. This may be due to a network configuration issue or temporary outage	**'app'** - The application name. **'attempts'** - The number of attempts tried.
	Warning. The application is abandoning the update.	
UPDATER.UPDATERD.MANIFEST_ FAILED_ALERT	The updater has been unable to communicate with the update server for at least $threshold.	**'threshold'** - Human readable threshold string.
	Warning. Failed to acquire a server manifest.	
UPDATER.UPDATERD.RELEASE_NOTIFICATION	$mail_text	**'mail_text'** - The notification text. **'notification_subject'** - The notification text.
	Warning. Release notification.	
UPDATER.UPDATERD.UPDATE_FAILED	Unknown error occured: $traceback	**'traceback'** - The traceback.
	Critical. Failed to run an update.	

Outbreak Filter Alerts

Table 15-10 contains a list of the various Outbreak Filter alerts that can be generated by AsyncOS, including a description of the alert and the alert severity. Please note that Outbreak Filters can also be referenced in system alerts for quarantines (the Outbreak quarantine, specifically).

Table 15-10 Listing of Possible Outbreak Filter Alerts

Alert Name	Message and Description	Parameters
VOF.GTL_THRESHOLD_ALERT	Cisco IronPort Outbreak Filters Rule Update Alert:$text All rules last updated at: $time on $date.	'**text**' - Update alert text. '**time**' - Time of last update. '**date**' - Date of last update.
	Information. Sent when the Outbreak Filters threshold has changed.	
AS.UPDATE_FAILURE	$engine update unsuccessful. This may be due to transient network or DNS issues, HTTP proxy configuration causing update transmission errors or unavailability of downloads.ironport.com. The specific error on the appliance for this failure is: $error	'**engine**' - The engine that failed to update. '**error**' - The error that happened.
	Warning. Sent when the anti-spam engine or CASE rules fail to update.	

Clustering Alerts

Table 15-10 contains a list of the various clustering alerts that can be generated by AsyncOS, including a description of the alert and the alert severity.

Table 15-11 Listing of Possible Clustering Alerts

Alert Name	Message and Description	Parameters
CLUSTER.CC_ERROR.AUTH_ERROR	Error connecting to cluster machine $name at IP $ip - $error - whyerror:=Machine does not appear to be in the cluster	'**name**' - The hostname and/or serial number of the machine. '**ip**' - The IP of the remote host. '**why**' - Detailed text about the error.
	Critical. Sent when there was an authentication error. This can occur if a machine is not a member of the cluster.	
CLUSTER.CC_ERROR.DROPPED	Error connecting to cluster machine $name at IP $ip - $error - whyerror:=Existing connection dropped	'**name**' - The hostname and/or serial number of the machine. '**ip**' - The IP of the remote host. '**why**' - Detailed text about the error.
	Warning. Sent when the connection to the cluster was dropped.	

Table 15-11 Listing of Possible Clustering Alerts (Continued)

Alert Name	Message and Description	Parameters
CLUSTER.CC_ERROR.FAILED	Error connecting to cluster machine $name at IP $ip - $error - whyerror:=Connection failure Warning. Sent when the connection to the cluster failed.	'**name**' - The hostname and/or serial number of the machine. '**ip**' - The IP of the remote host. '**why**' - Detailed text about the error.
CLUSTER.CC_ERROR.FORWARD_ FAILED	Error connecting to cluster machine $name at IP $ip - $error - whyerror:=Message forward failed, no upstream connection Critical. Sent when the appliance was unable to forward data to a machine in the cluster.	'**name**' - The hostname and/or serial number of the machine. '**ip**' - The IP of the remote host. '**why**' - Detailed text about the error.
CLUSTER.CC_ERROR.NOROUTE	Error connecting to cluster machine $name at IP $ip - $error - whyerror:=No route found Critical. Sent when the machine was unable to obtain a route to another machine in the cluster.	'**name**' - The hostname and/or serial number of the machine. '**ip**' - The IP of the remote host. '**why**' - Detailed text about the error.
CLUSTER.CC_ERROR.SSH_KEY	Error connecting to cluster machine $name at IP $ip - $error - whyerror:=Invalid host key Critical. Sent when there was an invalid SSH host key.	'**name**' - The hostname and/or serial number of the machine. '**ip**' - The IP of the remote host. '**why**' - Detailed text about the error.
CLUSTER.CC_ERROR.TIMEOUT	Error connecting to cluster machine $name at IP $ip - $error - whyerror:=Operation timed out Warning. Sent when the specified operation timed out.	'**name**' - The hostname and/or serial number of the machine. '**ip**' - The IP of the remote host. '**why**' - Detailed text about the error.

Table 15-11 Listing of Possible Clustering Alerts (Continued)

Alert Name	Message and Description	Parameters
CLUSTER.CC_ERROR_NOIP	Error connecting to cluster machine $name - $error - $why	'**name**' - The hostname and/or serial number of the machine. '**why**' - Detailed text about the error.
	Critical. Sent when the appliance could not obtain a valid IP address for another machine in the cluster.	
CLUSTER.CC_ERROR_NOIP. AUTH_ERROR	Error connecting to cluster machine $name - $error - whyerror:=Machine does not appear to be in the cluster	'**name**' - The hostname and/or serial number of the machine. '**why**' - Detailed text about the error.
	Critical. Sent when there was an authentication error connecting to a machine in a cluster. This can occur if a machine is not a member of the cluster.	
CLUSTER.CC_ERROR_ NOIP.DROPPED	Error connecting to cluster machine $name - $error - whyerror:=Existing connection dropped	'**name**' - The hostname and/or serial number of the machine. '**why**' - Detailed text about the error.
	Warning. Sent when the machine was unable to obtain a valid IP address for another machine in the cluster and the connection to the cluster was dropped.	
CLUSTER.CC_ERROR_ NOIP.FAILED	Error connecting to cluster machine $name - $error - whyerror:=Connection failure	'**name**' - The hostname and/or serial number of the machine. '**why**' - Detailed text about the error.
	Warning. Sent when there was an unknown connection failure and the machine was unable to obtain a valid IP address for another machine in the cluster.	

Table 15-11 Listing of Possible Clustering Alerts (Continued)

Alert Name	Message and Description	Parameters
CLUSTER.CC_ERROR_ NOIP.FORWARD_FAILED	Error connecting to cluster machine $name - $error - whyerror:=Message forward failed, no upstream connection	'**name**' - The hostname and/or serial number of the machine. '**why**' - Detailed text about the error.
	Critical. Sent when the machine was unable to obtain a valid IP address for another machine in the cluster and the appliance was unable to forward data to the machine.	
CLUSTER.CC_ERROR_ NOIP.NOROUTE	Error connecting to cluster machine $name - $error - whyerror:=No route found	'**name**' - The hostname and/or serial number of the machine. '**why**' - Detailed text about the error.
	Critical. Sent when the machine was unable to obtain a valid IP address for another machine in the cluster and it was unable to obtain a route to the machine.	
CLUSTER.CC_ERROR_ NOIP.SSH_KEY	Error connecting to cluster machine $name - $error - whyerror:=Invalid host key	'**name**' - The hostname and/or serial number of the machine. '**why**' - Detailed text about the error.
	Critical. Sent when the machine was unable to obtain a valid IP address for another machine in the cluster and was unable to obtain a valid SSH host key.	
CLUSTER.CC_ERROR_ NOIP.TIMEOUT	Error connecting to cluster machine $name - $error - whyerror:=Operation timed out	'**name**' - The hostname and/or serial number of the machine. '**why**' - Detailed text about the error.
	Warning. Sent when the machine was unable to obtain a valid IP address for another machine in the cluster and the specified operation timed out.	
CLUSTER.SYNC.PUSH_ALERT	Overwriting $sections on machine $name	'**name**' - The hostname and/or serial number of the machine. '**sections**' - List of cluster sections being sent.
	Critical. Sent when configuration data has gotten out of sync and has been sent to a remote host.	

Changing Network Settings

This section describes the features used to configure the network operation of the Cisco IronPort appliance. These features give you direct access to the hostname, DNS, and routing settings that you configured via the System Setup Wizard (or the `systemsetup` command) in Using the System Setup Wizard, page 3-13.

The following features are described:

- `sethostname`
- DNS Configuration (GUI and via the `dnsconfig` command)
- Routing Configuration (GUI and via the `routeconfig` and `setgateway` commands)
- `dnsflush`
- Password
- Network Access
- Login Banner

Changing the System Hostname

The hostname is used to identify the system at the CLI prompt. You must enter a fully-qualified hostname. The `sethostname` command sets the name of the Cisco IronPort appliance. The new hostname does not take effect until you issue the `commit` command.

The sethostname Command

```
oldname.example.com> sethostname


[oldname.example.com]> mail3.example.com


oldname.example.com>
```

For the hostname change to take effect, you must enter the `commit` command. After you have successfully committed the hostname change, the new name appears in the CLI prompt:

```
oldname.example.com> commit


Please enter some comments describing your changes:

[]> Changed System Hostname

Changes committed: Mon Jan 01 12:00:01 2003
```

The new hostname appears in the prompt as follows: `mail3.example.com>`

Configuring Domain Name System (DNS) Settings

You can configure the DNS settings for your Cisco IronPort appliance through the DNS page on the Network menu of the GUI, or via the `dnsconfig` command.

You can configure the following settings:

- whether to use the Internet's DNS servers or your own, and which specific server(s) to use
- which interface to use for DNS traffic
- the number of seconds to wait before timing out a reverse DNS lookup
- clear DNS cache

Specifying DNS Servers

Cisco IronPort AsyncOS can use the Internet root DNS servers, your own DNS servers, or the Internet root DNS servers and authoritative DNS servers you specify. When using the Internet root servers, you may specify alternate servers to use for specific domains. Since an alternate DNS server applies to a single domain, it must be authoritative (provide definitive DNS records) for that domain.

AsyncOS supports "splitting" DNS servers when not using the Internet's DNS servers. If you are using your own internal server, you can also specify exception domains and associated DNS servers.

When setting up "split DNS," you should set up the in-addr.arpa (PTR) entries as well. So, for example, if you want to redirect ".eng" queries to the nameserver 1.2.3.4 and all the .eng entries are in the 172.16 network, then you should specify "eng,16.172.in-addr.arpa" as the domains in the split DNS configuration.

Multiple Entries and Priority

For each DNS server you enter, you can specify a numeric priority. AsyncOS will attempt to use the DNS server with the priority closest to 0. If that DNS server is not responding AsyncOS will attempt to use the server at the next priority. If you specify multiple entries for DNS servers with the same priority, the system randomizes the list of DNS servers at that priority every time it performs a query. The system then waits a short amount of time for the first query to expire or "time out" and then a slightly longer amount of time for the second, etc. The amount of time depends on the exact total number of DNS servers and priorities that have been configured. The timeout length is the same for all IP addresses at any particular priority. The first priority gets the shortest timeout, each subsequent priority gets a longer timeout. Further, the timeout period is roughly 60 seconds. If you have one priority, the timeout for each server at that priority will be 60 seconds. If you have two priorities, the timeout for each server at the first priority will be 15 seconds, and each server at the second priority will be 45 seconds. For three priorities, the timeouts are 5, 10, 45.

For example, suppose you configure four DNS servers, with two of them at priority 0, one at priority 1, and one at priority 2:

Table 15-12 Example of DNS Servers, Priorities, and Timeout Intervals

Priority	Server(s)	Timeout (seconds)
0	1.2.3.4, 1.2.3.5	5, 5
1	1.2.3.6	10
2	1.2.3.7	45

AsyncOS will randomly choose between the two servers at priority 0. If one of the priority 0 servers is down, the other will be used. If both of the priority 0 servers are down, the priority 1 server (1.2.3.6) is used, and then, finally, the priority 2 (1.2.3.7) server.

The timeout period is the same for both priority 0 servers, longer for the priority 1 server, and longer still for the priority 2 server.

Using the Internet Root Servers

The Cisco IronPort AsyncOS DNS resolver is designed to accommodate the large number of simultaneous DNS connections required for high-performance email delivery.

Note If you choose to set the default DNS server to something other than the Internet root servers, that server must be able to recursively resolve queries for domains for which it is not an authoritative server.

Reverse DNS Lookup Timeout

The Cisco IronPort appliance attempts to perform a "double DNS lookup" on all remote hosts connecting to a listener for the purposes of sending or receiving email. [That is: the system acquires and verifies the validity of the remote host's IP address by performing a double DNS lookup. This consists of a reverse DNS (PTR) lookup on the IP address of the connecting host, followed by a forward DNS (A) lookup on the results of the PTR lookup. The system then checks that the results of the A lookup match the results of the PTR lookup. If the results do not match, or if an A record does not exist, the system only uses the IP address to match entries in the Host Access Table (HAT).] This particular timeout period applies only to this lookup and is not related to the general DNS timeout discussed in Multiple Entries and Priority, page 15-39.

The default value is 20 seconds. You can disable the reverse DNS lookup timeout globally across all listeners by entering '0' as the number of seconds.

If the value is set to 0 seconds, the reverse DNS lookup is not attempted, and instead the standard timeout response is returned immediately. This also prevents the appliance from delivering mail to domains that require TLS-verified connections if the receiving host's certificate has a common name (CN) that maps to the host's IP lookup.

DNS Alert

Occasionally, an alert may be generated with the message "Failed to bootstrap the DNS cache" when an appliance is rebooted. The messages means that the system was unable to contact its primary DNS servers, which can happen at boot time if the DNS subsystem comes online before network connectivity is established. If this message appears at other times, it could indicate network issues or that the DNS configuration is not pointing to a valid server.

Clearing the DNS Cache

The Clear Cache button from the GUI, or the `dnsflush` command (for more information about the `dnsflush` command, see the *Cisco IronPort AsyncOS CLI Reference Guide*), clears all information in the DNS cache. You may choose to use this feature when changes have been made to your local DNS system. The command takes place immediately and may cause a temporary performance degradation while the cache is repopulated.

Configuring DNS Settings via the Graphical User Interface

Step 1 Select **Network > DNS**.

Step 2 Click **Edit Settings**. The Edit DNS page is displayed:

Figure 15-12 The Edit DNS Page

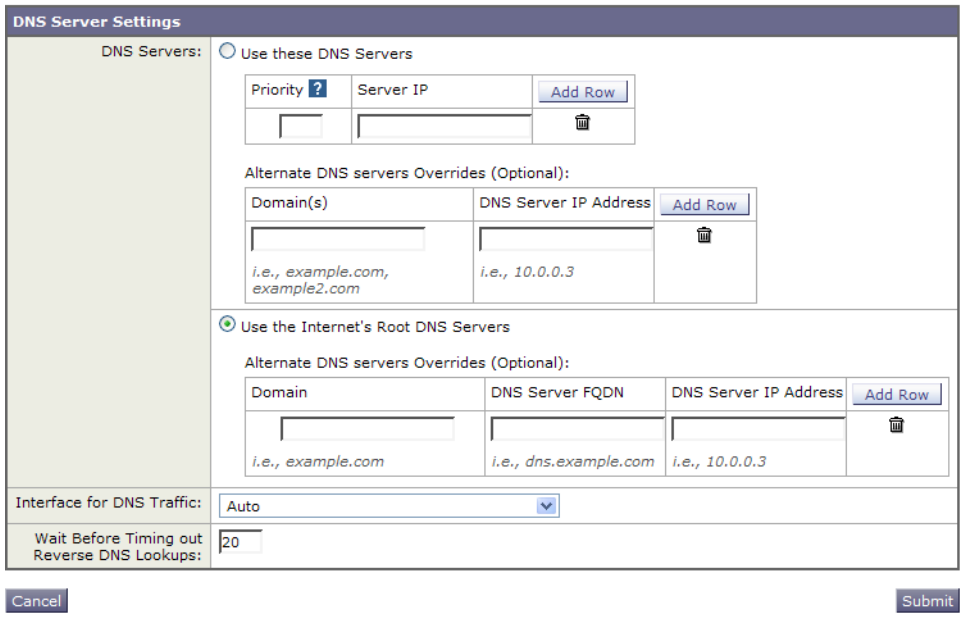

Step 3 Select whether to use the Internet's root DNS servers or your own internal DNS server or the Internet's root DNS servers and specify alternate DNS servers.

Step 4 If you want to use your own DNS server(s) enter the server ID and click **Add Row**. Repeat this for each server. When entering your own DNS servers, specify a priority as well. For more information, see Specifying DNS Servers, page 15-39.

Step 5 If you want to specify alternate DNS servers for certain domains, enter the domain and the alternate DNS server IP address. Click **Add Row** to add additional domains.

Note You can enter multiple domains for a single DNS server by using commas to separate domain names. You can also enter multiple DNS servers by using commas to separate IP addresses.

Step 6 Choose an interface for DNS traffic.

Step 7 Enter the number of seconds to wait before cancelling a reverse DNS lookup.

Step 8 You can also clear the DNS cache by clicking **Clear Cache**.

Step 9 Submit and commit your changes.

Configuring TCP/IP Traffic Routes

Some network environments require the use of traffic routes other than the standard default gateway. You can manage static routes via the GUI through the Routing page on the Network tab, or the CLI, via the `routeconfig` command.

Managing Static Routes (GUI)

You can create, edit, or delete static routes via the Routing page on the Network tab. The Email Security appliance can use both Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) static routes, which you can create and manage separately on the Routing page. You can also modify the default IPv4 and IPv6 gateways from this page.

Adding Static Routes

Step 1 Click **Add Route** for the type of static route you want to create on the Routing page. The Add Static Route page is displayed. The following figure shows the IPv6 Static Route Settings.

Figure 15-13 Adding a Static Route

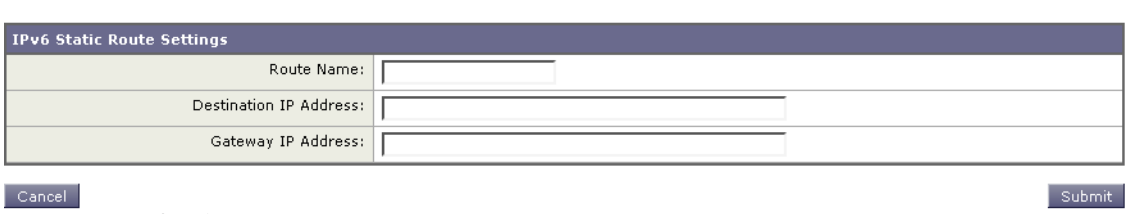

Step 2 Enter a name for the route.

Step 3 Enter the destination IP address.

Step 4 Enter the Gateway IP address.

Step 5 Submit and commit your changes.

Deleting Static Routes

Step 1 Click the trash can icon corresponding to the static route name in the Static Routes listing.

Step 2 Confirm the deletion by clicking **Delete** in the warning dialog that appears.

Step 3 Commit your changes.

Editing Static Routes

Step 1 Click the name of the route in the Static Route listing. The Edit Static Route page is displayed.

Step 2 Make changes to the route.

Step 3 Commit your changes.

Modifying the Default Gateway

Step 1 Click Default Route in the route listing for the Internet Protocol version you want to modify on the Routing page. The Edit Default Route page is displayed. The following figure shows the Edit Default Route page for the IPv6 default gateway.

Figure 15-14 Editing the Default Gateway

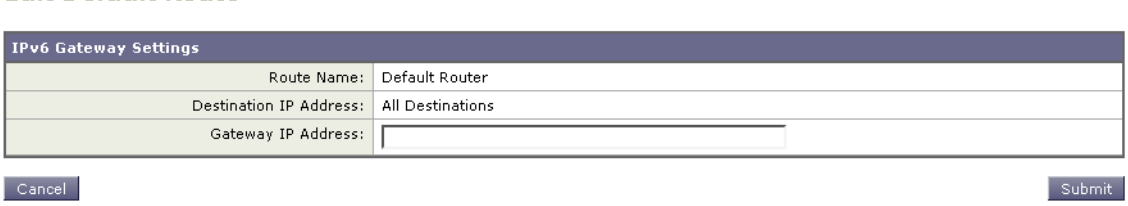

Step 2 Change the Gateway IP address.

Step 3 Submit and commit your changes.

Configuring the Default Gateway

You can configure the default gateway via the GUI though the Static Routes page on the Network menu (see Modifying the Default Gateway, page 15-43) or via the `setgateway` command in the CLI.

Changing the admin User's Password

The password for the admin user can be changed via the GUI or the CLI.

To change the password via the GUI, use the Users page available via the System Administration tab. For more information, see the section on managing users in "Common Administrative Tasks" in the *Cisco IronPort AsyncOS for Email Daily Management Guide*.

To change the password for the `admin` user via the CLI, use the `password` command. Passwords must be six characters or longer. The `password` command requires you to enter the old password for security.

Note Changes to the password take effect immediately and do not require you to execute the `commit` command.

Configuring Access to the Email Security Appliance

AsyncOS provides administrators controls to manage users' access to the Email Security appliance, including a timeout for Web UI session and an access list that specifies the IP addresses from which users and your organization's proxy servers can access the appliance.

Configuring IP-Based Network Access

You can control from which IP addresses users access the Email Security appliance by creating access lists for users who connect directly to the appliance and users who connect through a reverse proxy, if your organization uses reverse proxies for remote users.

Direct Connections

You can specify the IP addresses, subnets, or CIDR addresses for machines that can connect to the Email Security appliance. Users can access the appliance from any machine with IP address from the access list. Users attempting to connect to the appliance from an address not included in the list are denied access.

Connecting Through a Proxy

If your organization's network uses reverse proxy servers between remote users' machines and the Email Security appliance, AsyncOS allows you create an access list with the IP addresses of the proxies that can connect to the appliance.

Even when using a reverse proxy, AsyncOS still validates the IP address of the remote user's machine against a list of IP addresses allowed for user connections. To send the remote user's IP address to the Email Security appliance, the proxy needs to include the `x-forwarded-for` HTTP header in its connection request to the appliance.

The `x-forwarded-for` header is a non-RFC standard HTTP header with the following format:

```
x-forwarded-for: client-ip, proxy1, proxy2,... CRLF.
```

The value for this header is a comma-separated list of IP addresses with the left-most address being the address of the remote user's machine, followed by the addresses of each successive proxy that forwarded the connection request. (The header name is configurable.) The Email Security appliance matches the remote user's IP address from the header and the connecting proxy's IP address against the allowed user and proxy IP addresses in the access list.

Note AsyncOS supports only IPv4 addresses in the `x-forwarded-for` header.

Creating the Access List

You can create the network access list either via the Network Access page in the GUI or the `adminaccessconfig > ipaccess` CLI command. Figure 15-15 shows the Network Access page with a list of user IP addresses that are allowed to connect directly to the Email Security appliance.

Figure 15-15 Network Access Settings

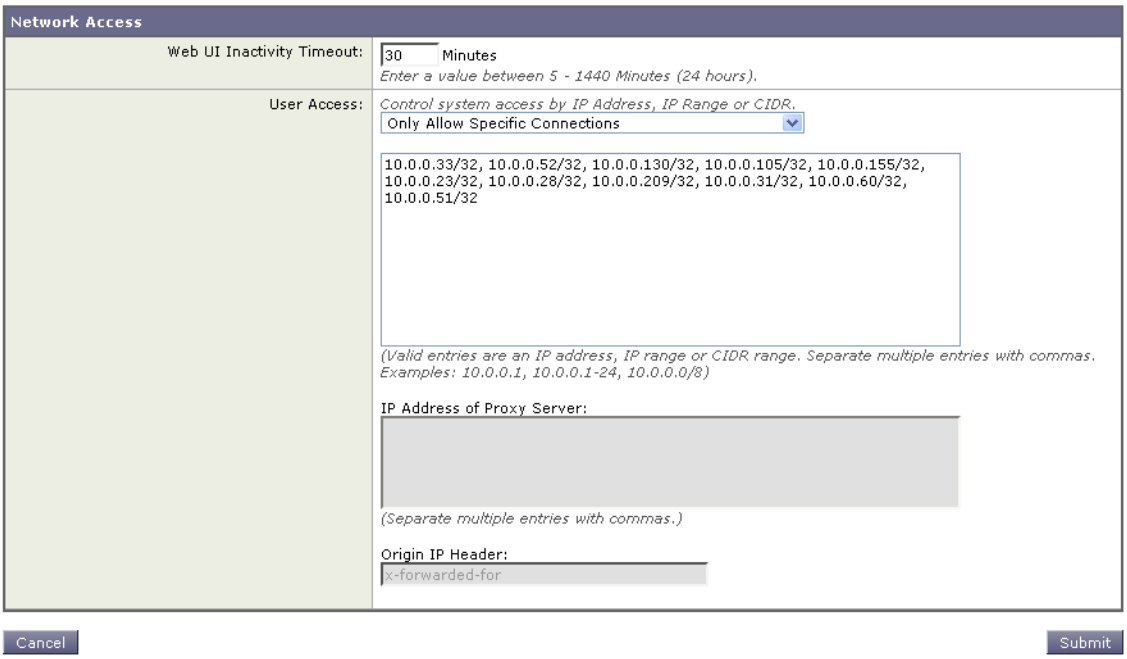

AsyncOS offers four different modes of control for the access list:

- **Allow All.** This mode allows all connections to the appliance. This is the default mode of operation.
- **Only Allow Specific Connections.** This mode allows a user to connection to the appliance if the user's IP address matches the IP addresses, IP ranges, or CIDR ranges included in the access list.
- **Only Allow Specific Connections Through Proxy.** This mode allows a user to connect to the appliance through a reverse proxy if the following conditions are met:
  - The connecting proxy's IP address is included in the access list's IP Address of Proxy Server field.
  - The proxy includes the `x-forwarded-header` HTTP header in its connection request.
  - The value of `x-forwarded-header` is not empty.
  - The remote user's IP address is included in `x-forwarded-header` and it matches the IP addresses, IP ranges, or CIDR ranges defined for users in the access list.
- **Only Allow Specific Connections Directly or Through Proxy.** This mode allows users to connect through a reverse proxy or directly to the appliance if their IP address matches the IP addresses, IP ranges, or CIDR ranges included in the access list. The conditions for connecting through a proxy are the same as in the Only Allow Specific Connections Through Proxy mode.

Please be aware that you may lose access to the appliance after submitting and committing your changes if one of the following conditions is true:

- If you select **Only Allow Specific Connections** and do not include the IP address of your current machine in the list.
- If you select **Only Allow Specific Connections Through Proxy** and the IP address of the proxy currently connected to the appliance is not in the proxy list and the value of the Origin IP header is not in the list of allowed IP addresses.
- If you select **Only Allow Specific Connections Directly or Through Proxy** and

- the value of the Origin IP header is not in the list of allowed IP addresses

  OR

- the value of the Origin IP header is not in the list of allowed IP Addresses and the IP address of the proxy connected to the appliance is not in the list of allowed proxies.

Step 1 Use the System Administration > Network Access page.

Step 2 Click **Edit Settings**.

Step 3 Select the mode of control for the access list.

Step 4 Enter the IP addresses from which users will be allowed to connect to the appliance.

You can enter an IP address, IP address range or CIDR range. Use commas to separate multiple entries.

Step 5 If connecting through a proxy is allowed, enter the following information:

- The IP addresses of the proxies allowed to connect to the appliance. Use commas to separate multiple entries.
- The name of the origin IP header that the proxy sends to the appliance, which contains the IP addresses of the remote user's machine and the proxy servers that forwarded the request. By default, the name of the header is `x-forwarded-for`.

Step 6 Submit and commit your changes.

Configuring the Web UI Session Timeout

You can specify how long a user can be logged into the Email Security appliance's Web UI before AsyncOS logs the user out due to inactivity. This Web UI session timeout applies to all users, including admin, and it is used for both HTTP and HTTPS sessions.

Once AsyncOS logs a user out, the appliance redirects the user's web browser to login page.

Note The Web UI Session Timeout does not apply to Cisco IronPort Spam Quarantine sessions, which have a 30 minute timeout that cannot be configured.

Figure 15-16 Web UI Inactivity Timeout

Step 1 Use the System Administration > Network Access page.

Step 2 Click **Edit Settings**.

Step 3 Enter the number of minutes users can be inactive before being logged out. You can define a timeout period between 5 and 1440 minutes.

Step 4 Submit and commit your changes.

Adding a Login Banner

You can configure the Email Security appliance to display a message called a "login banner" when a user attempts to log into the appliance through SSH, Telnet, FTP, or Web UI. The login banner is customizable text that appears above the login prompt in the CLI and to the right of the login prompt in the GUI. You can use the login banner to display internal security information or best practice instructions for the appliance. For example, you can create a simple note that saying that unauthorized use of the appliance is prohibited or a detailed warning concerning the organization's right to review changes made by the user to the appliance.

Use the `adminaccessconfig > banner` command in the CLI to create the login banner. The maximum length of the login banner is 2000 characters to fit 80x25 consoles. A login banner can be imported from a file in the `/data/pub/configuration` directory on the appliance. After creating the banner, commit your changes.

Figure 15-17 shows a login banner displayed on the Web UI login screen.

Figure 15-17 Web UI Login Screen with Banner

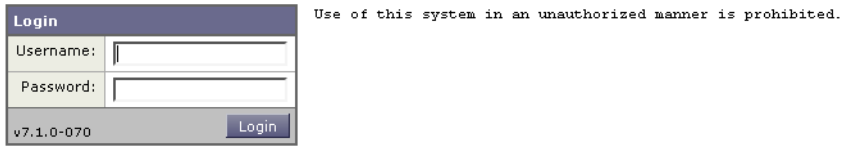

System Time

To set the System Time on your Cisco IronPort appliance, set the Time Zone used, or select an NTP server and query interface, use the Time Zone or Time Settings page from the System Administration menu in the GUI or use the following commands in the CLI: `ntpconfig`, `settime`, and `settz`.

You can also verify the time zone files used by AsyncOS on the System Administration > Time Settings page or using the `tzupdate` CLI command.

Selecting a Time Zone

The Time Zone page (available via the System Administration menu in the GUI) displays the time zone for your Cisco IronPort appliance. You can select a specific time zone or GMT offset.

Step 1 Click **Edit Settings** on the System Administration > Time Zone page. The Edit Time Zone page is displayed:

Figure 15-18 The Time Zone Page

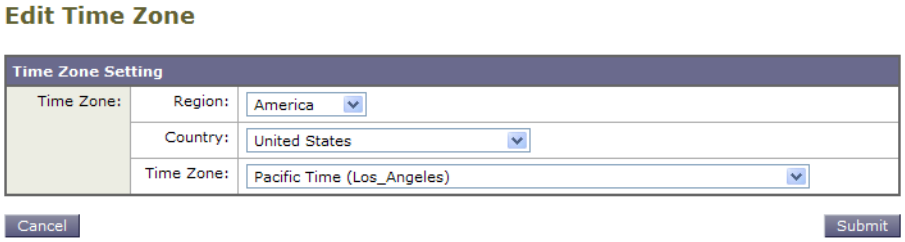

Step 2 Select a Region, country, and time zone from the pull-down menus.

Step 3 Submit and commit your changes.

Selecting a GMT Offset

Step 1 Click **Edit Settings** on the System Administration > Time Zone page. The Edit Time Zone page is displayed.

Step 2 Select GMT Offset from the list of regions. The Time Zone Setting page is updated:

Figure 15-19 The Time Zone Page

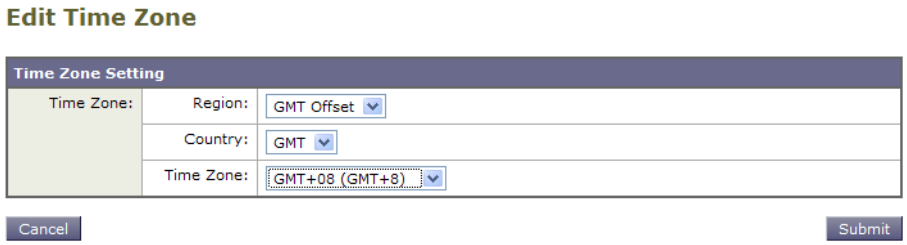

Step 3 Select an offset in the Time Zone list. The offset refers to the amount of hours that must be added/subtracted in order to reach GMT (the Prime Meridian). Hours preceded by a minus sign ("-") are east of the Prime Meridian. A plus sign ("+") indicates west of the Prime Meridian.

Step 4 Submit and commit your changes.

Editing Time Settings

To edit the time settings for your Cisco IronPort appliance, click the **Edit Settings** button on the System Administration > Time Settings page. The Edit Time Settings page is displayed:

Figure 15-20 The Edit Time Settings Page

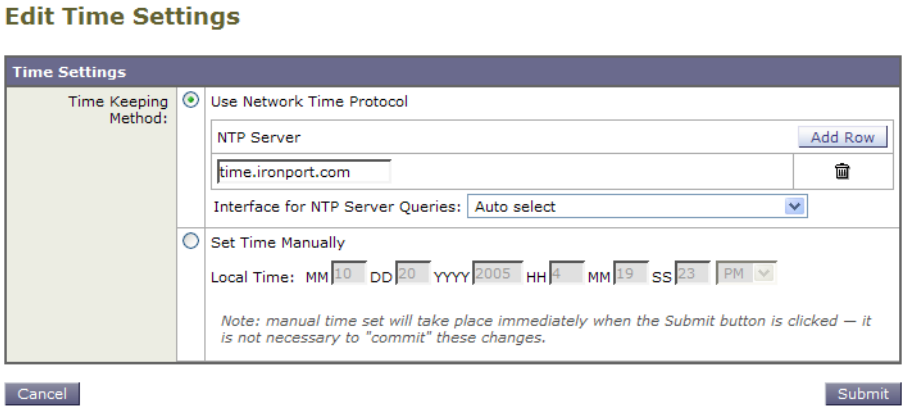

Editing the Network Time Protocol (NTP) Configuration (Time Keeping Method)

Step 1 Click **Edit Settings** on the System Administration > Time Settings page. The Edit Time Settings page is displayed.

Step 2 In the Time Keeping Method section, select Use Network Time Protocol.

Step 3 Enter an NTP server address and click **Add Row**. You can add multiple NTP servers.

Step 4 To delete an NTP server from the list, click the trash can icon for that server.

Step 5 Select an interface for NTP queries. This is the IP address from which NTP queries should originate.

Step 6 Submit and commit your changes.

Setting System Time (not using NTP Server)

Step 1 Click **Edit Settings** on the System Administration > Time Settings page. The Edit Time Settings page is displayed.

Step 2 In the Time Keeping Method section, select Set Time Manually.

Step 3 Enter the month, day, year, hour, minutes, and seconds.

Step 4 Select A.M or P.M.

Step 5 Submit and commit your changes.

CHAPTER 16

Enabling Your C350D Appliance

The C350D/C360D/C370D appliance is a special model of the Cisco IronPort appliances, specifically designed for outbound email delivery. This chapter discusses the various features of and modifications to the AsyncOS operating system specific to the C350D appliance. Note that in this chapter, the C350D, C360D, and C370D, Cappliances are interchangeable. The remainder of the text in this chapter refers only to the C350D; however, all information discussed is also applicable to the C370D and C360D appliances.

Overview: The C350D Appliance

The C350D appliance is a C350/360/370 appliance with a feature key for AsyncOS modifications designed and optimized for outbound delivery of mail. The C350D appliance includes dramatically enhanced performance intended to meet the specific needs of outbound customer messaging.

Additional Features for the C350D

As part of the optimization for message delivery, the C350D appliance contains some additional features not found in the standard Cisco IronPort appliances.

Additional features:

- 256 Virtual Gateway Addresses - The Cisco IronPort Virtual Gateway technology allows you to configure enterprise mail gateways for all domains you host — with distinct IP addresses, hostname and domains — and create separate corporate email policy enforcement and anti-spam strategies for those domains, while hosted within the same physical appliance. For more information, see “Customizing Listeners” in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.
- IronPort Mail Merge (IPMM) - IronPort Mail Merge (IPMM) removes the burden of generating individual personalized messages from customer systems. By removing the need to generate thousands of individual messages and transmit them between message generating systems and the email gateway, users benefit from the decreased load on their systems and increased throughput of email delivery. For more information, see IronPort Mail Merge (IPMM), page 16-4.

- Resource-conserving bounce setting - The C350D appliance allows you to configure the system to detect potential blocked destinations and bounce all messages bound for that destination. For more information, see Configuring Resource-Conserving Bounce Settings, page 16-4.
- Software based performance enhancement - The C350D appliance includes a software module that dramatically enhances the outbound delivery performance.

Features Disabled in the C350D

Your C350D appliance contains several modifications to the AsyncOS operating system. Some features of the standard C- and X-Series appliances are not applicable to outbound email delivery and to improve system performance have been disabled. These modifications and differences are discussed below.

Non-Applicable Features:

- Cisco IronPort anti-spam scanning and on or off box spam quarantining — Because anti-spam scanning pertains mostly to incoming mail, the Cisco IronPort Anti-Spam scanning engine is disabled. Chapter 9 is, therefore, not applicable.
- Outbreak Filters — Because Cisco IronPort's Outbreak Filters feature is used to quarantine incoming mail, this feature has been disabled on the C350D. Chapter 11 is, therefore, not applicable.
- SenderBase Network Participation capabilities — Because SenderBase Network Participation reports information about incoming mail, this feature has been disabled on the C350D appliance. Chapters 8 and 12 are, therefore, not applicable.
- Reporting — Reporting is limited. Some reports are not available, and the reporting that does occur is set to run at a very limited level due to the performance issues.
- RSA Data Loss Prevention — RSA DLP scanning for outgoing messages has been disabled on C350D appliances.
- The totals shown in the Email Security Monitor Overview report for 350D appliances may erroneously include spam and suspect spam counts although these features are disabled for the C350D appliances.

AsyncOS Features Applicable to the C350D

The C350D appliance incorporates most of the latest AsyncOS features, many of which are of interest to C350D users. Table 16-1 lists some of these features:

Table 16-1 AsyncOS Features Included in the C350D Appliance

Feature	More Information
Domain Key signing	DKIM/Domain Keys is a method for verifying authenticity of email based on a signing key used by the sender. See the "Email Authentication" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.
Centralized management	See the "Centralized Management" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Table 16-1 AsyncOS Features Included in the C350D Appliance (Continued)

Feature	More Information
Delivery throttling	For each domain, you can assign a maximum number of connections and recipients that will never be exceeded by the system in a given time period. This "good neighbor" table is defined through the `destconfig` command. For more information, see the section on Controlling Email Delivery in "Configuring Routing and Delivery Features" the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.
Bounce Verification	Verify the authenticity of bounce messages. See the section on Cisco IronPort Bounce Verification in the "Configuring Routing and Delivery Features" chapter of the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.
Delegated administration	See information on adding users in the "Common Administrative Tasks" chapter of the *Cisco IronPort AsyncOS for Email Daily Management Guide*.
Trace (debug)	See Debugging Mail Flow Using Test Messages: Trace, page -446.
VLAN, NIC-pairing	See the "Advanced Network Configuration" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.
Optional Anti-virus engine	You can add optional anti-virus scanning to ensure the integrity of your outbound messages. See Anti-Virus Scanning, page 8-1.

Configuring the C350D Appliance

Step 1 Apply the provided feature key. You will need to apply the key to your C350D Cisco IronPort Email Security appliance *prior to running the system setup wizard* (prior to configuring the appliance). Apply the key via the System Administration > Feature Key page or by issuing the `featurekey` command in the CLI.

Note The preceding feature keys include a sample 30 day Sophos or McAfee Anti-Virus license you can use to test anti-virus scanning on outbound mail.

Step 2 Reboot the appliance.

Step 3 Run the system setup wizard (GUI or CLI) and configure your appliance.

Please keep in mind that the Cisco IronPort C350D appliance does not include anti-spam scanning or the Outbreak Filters feature. (Please ignore these chapters in the Configuration Guide.)

Note In a clustered environment, you cannot combine C350D appliances with AsyncOS appliances that are not configured with the delivery performance package.

Configuring Resource-Conserving Bounce Settings

Once the C350D appliance is configured, you can configure the system to detect potential delivery problems and bounce all messages for a destination.

Note Using this setting will bounce all messages in the queue for a destination domain that is deemed undeliverable. You will need to re-send the message once the delivery issues have been resolved.

Example of Enabling Resource-Conserving Bounce Settings

```
mail3.example.com> bounceconfig


Choose the operation you want to perform:

- NEW - Create a new profile.

- EDIT - Modify a profile.

- DELETE - Remove a profile.

- SETUP - Configure global bounce settings.

[]> setup


Do you want to bounce all enqueued messages bound for a domain if the host is down? [N]>
y
```

When using this feature, a host is considered "down" after at least 10 consecutive connection attempts fail. AsyncOS scans for down hosts every 15 minutes, so it is possible that more than 10 attempts will be made before the queue is cleared.

IronPort Mail Merge (IPMM)

Note IronPort Mail Merge is only available on the IronPort C350D appliance.

Overview

IronPort Mail Merge removes the burden of generating individual personalized messages from customer systems. By removing the need to generate thousands of individual messages and transmit them between message generating systems and the email gateway, users benefit from the decreased load on their systems and increased throughput of email delivery.

With IPMM, a single message body is created with variables representing locations in the message to be replaced for personalization. For each individual message recipient, only the recipient email address and the variable substitutions need to be transmitted to the email gateway. In addition, IPMM can be used to send certain recipients specific "parts" of the message body, while excluding certain parts from others recipients. (For example, suppose you needed to include a different copyright statements at the end of your messages to recipients in two different countries.)

Benefits

Using the Mail Merge function of the Cisco IronPort C350D appliance has many benefits:

- Ease of use for the mail administrator. The complexities of creating personalized messages for each recipient are removed, as IPMM provides variable substitution and an abstracted interface in many common languages.
- Reduced load on message generation systems. By requiring one copy of the message body and a table of required substitutions, most of the message generation "work" is off-loaded from message generation systems and moved to the Cisco IronPort C350D appliance.
- Increased delivery throughput. By reducing the resources necessary to accept and queue thousands of incoming messages, the Cisco IronPort appliance can significantly increase out-bound delivery performance.
- Queue storage efficiency. By storing less information for each message recipient, users can achieve orders-of- magnitude, better use of queue storage on the C350D appliance.

Using the Mail Merge

SMTP Injection

IPMM extends SMTP as the transport protocol. There is no special configuration that needs to be made to the Cisco IronPort C350D appliance. (By default, IPMM can be enabled for private listeners and disabled for public listeners on the Cisco IronPort C350D Email Security appliance.) However, if you are not currently using SMTP as your injection protocol, you must create a new private listener that utilizes SMTP through the Cisco IronPort C350D appliance interface.

Refer to the "Customizing Listeners" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide* for more information on configuring listeners. Use the `setipmm` subcommand of `listenerconfig` to enable IPMM on the injector.

IPMM modifies SMTP by altering two commands — `MAIL FROM` and `DATA` — and adding another: `XDFN`. The `MAIL FROM` command is replaced with `XMRG FROM` and, the `DATA` command is replaced with `XPRT`.

To generate a Mail Merge message, the commands used to generate the message need to be issued in a particular sequence.

1. The initial EHLO statement, identifying the sending host.
2. Each message starts with an XMRG FROM: statement, indicating the sender address.
3. Each recipient is then defined:
    - One or more XDFN variable allocation statements are made, including defining the parts (XDFN *PART=1,2,3…), and any other recipient specific variables.

- The recipient email address is defined with the RCPT TO: statement. Any variable allocations prior to the RCPT TO:, but after the prior XMRG FROM, or RCPT TO command, will be mapped to this recipient email address.

4. Each part is defined using the XPRT n command, with each part terminated by a period (.) character similar to the DATA command. The last part is defined by the XPRT n LAST command.

Variable Substitution

Any part of the message body, including message headers, can contain variables for substitution. Variables can appear in HTML messages, as well. Variables are user-defined and must begin with the ampersand (`&`) character and end with the semi-colon character (`;`). Variable names beginning with an asterisk (`*`) are reserved and cannot be used.

Reserved Variables

IPMM contains five special "reserved" variables that are predefined.

Table 16-2 IPMM: Reserved Variables

`*FROM`	The reserved variable `*FROM` is derived from the "Envelope From" parameter. The "Envelope From" parameter is set by the "XMRG FROM:" command.
`*TO`	The reserved variable `*TO` is derived from the envelope recipient value, as set by the "RCPT TO:" command.
`*PARTS`	The reserved variable `*PARTS` holds a comma separated list of parts. It is set prior to defining a recipient with the "RCPT TO:" and determines which of the "XPRT n" message body blocks a given user will receive.
`*DATE`	The reserved variable `*DATE` is replaced with the current date stamp.
`*DK`	The reserved variable `*DK` is used to specify a DomainKeys Signing profile (this profile must already exist in AsyncOS). For more information about creating DomainKeys Signing profiles, see the "Email Authentication" chapter in *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

For example, the following example message body (including headers) contains four distinct variables and five substitution locations that will be replaced in the final message. Note that the same variable may be used more than once in the message body. Also, the reserved variable `&*TO;` is used, which will be replaced with the recipient email address. This reserved variable does not need to be passed in as a separate variable. The variables in the example appear in bold.

Example Message #1

```
From: Mr.Spacely <spacely@sprockets.com>

To: &first_name;&last_name;&*TO;

Subject: Thanks for Being a Spacely Sprockets Customer


Dear &first_name;,


Thank you for purchasing a &color; sprocket.
```

This message needs only be injected once into the Cisco IronPort C350D appliance. For each recipient, the following additional information is required:

- A recipient email address
- Name-value pairs for the variable substitution

Part Assembly

Where SMTP uses a single `DATA` command for each message body, IPMM uses one or many `XPRT` commands to comprise a message. Parts are assembled based upon the order specified per-recipient. Each recipient can receive any or all of the message parts. Parts can be assembled in any order.

The special variable `*PARTS` holds a comma separated list of parts.

For example, the following example message contains two parts.

The first part contains the message headers and some of the message body. The second part contains an offer that can be variably included for specific customers.

Example Message #2, Part 1

```
From: Mr. Spacely <spacely@sprockets.com>

To: &first_name; &last_name; &*TO;

Subject: Thanks for Being a Spacely Sprockets Customer


Dear &first_name;,


Thank you for purchasing a &color; sprocket.
```

Example Message #2, Part 2

```
Please accept our offer for 10% off your next sprocket purchase.
```

The message parts need only be injected once into the Cisco IronPort C350D appliance. In this case, each recipient requires the following additional information:

- The ordered list of parts to be included in the final message
- A recipient email address
- Name value pairs for the variable substitution

IPMM and DomainKeys Signing

IPMM does support DomainKeys Signing. Use the `*DK` reserved variable to specify a DomainKeys profile. For example:

```
XDFN first_name="Jane" last_name="User" color="red" *PARTS=1,2 *DK=mass_mailing_1
```

In this example, “mail_mailing_1” is the name of a previously configured DomainKeys profile.

Command Descriptions

When a client injects IPMM messages to the listener, it uses extended SMTP with the following key commands.

XMRG FROM

Syntax:

XMRG FROM: <*sender email address*>

This command replaces the SMTP `MAIL FROM:` command and indicates that what follows is an IPMM message. An IPMM job is initiated with the `XMRG FROM:` command.

XDFN

Syntax:

```
XDFN <KEY=VALUE> [KEY=VALUE]
```

The XDFN command sets the per-recipient metadata. Note that key-value pairs can optionally be enclosed in angle brackets or square brackets.

*PARTS is a special reserved variable that indicates the index number as defined by the XPRT command (described below). The *PARTS variable is split as a comma-delimited list of integers. The integers match the body parts to be sent as defined by the XPRT commands. The other reserved variables are: *FROM, *TO, and *DATE.

XPRT

Syntax:

```
XPRT index_number LAST

Message

.
```

The XPRT command replaces the SMTP DATA command. The command accepts the transfer of the message part after the command is issued. The command is completed with a single period on a line followed by a return (which is the same way an SMTP DATA command is completed).

The special keyword **LAST** indicates the end of the mail merge job and must be used to specify the final part that will be injected.

After the LAST keyword is used, the message is queued, and delivery begins.

Notes on Defining Variables

- When you define variables with the XDFN command, note that the actual command line cannot exceed the physical limit of the system. In the case of the Cisco IronPort C350D appliance, this limit is 4 kilobytes per line. Other host systems may have lower thresholds. Use caution when defining multiple variables on very large lines.
- You can escape special characters using the forward slash "/" character when defining variables key-value pairs. This is useful if your message body contains HTML character entities that might be mistakenly replaced with variable definitions. (For example, the character entity ™ defines the HTML character entity for a trademark character. If you created the command XDFN trade=foo and then created a IPMM message containing the HTML character entity "™" the assembled message would contain the variable substitution ("foo") instead of the trademark character. The same concept is true for the ampersand character "&" which is sometimes used in URLs containing GET commands.

Example IPMM Conversation

The following is an example IPMM conversation of Example Message #2 (shown above). The message will be sent to two recipients in this example: "Jane User" and "Joe User."

In this example, the type in **bold** represents what you would type in a manual SMTP conversation with the Cisco IronPort C350D appliance, type in monospaced type represents the responses from the SMTP server, and *italic type* represents comments or variables.

A connection is established:

```
220 ESMTP

EHLO foo

250-ehlo responses from the injector enabled for IPMM
```

The conversation is started:

```
XMRG FROM:<user@domain.com> [Note: This replaces the MAIL FROM: SMTP command.]

250 OK
```

Variables and parts are set for each recipient:

```
XDFN first_name="Jane" last_name="User" color="red" *PARTS=1,2

[Note: This line defines three variables (first_name, last_name, and color) and then
uses the *PARTS reserved variable to define that the next recipient defined will receive
message parts numbers 1 and 2.]

250 OK

RCPT TO:<jane@example.com>

250 recipient <jane@example.com> ok

XDFN first_name="Joe" last_name="User" color="black" *PARTS=1

[Note: This line defines three variables (first_name, last_name, and color) and then
uses the *PARTS reserved variable to define that the next recipient defined will receive
message parts numbers 1 only.]

RCPT TO:<joe@example.com>

250 recipient <joe@example.com> ok
```

Next, part 1 is transmitted:

```
XPRT 1 [Note: This replaces the DATA SMTP command.]

354 OK, send part

From: Mr. Spacely <spacely@sprockets.com>

To: &first_name; &last_name; &*TO;

Subject: Thanks for Being a Spacely Sprockets Customer
```

```
&*DATE;

Dear &first_name;,


Thank you for purchasing a &color; sprocket.

.
```

And then part 2 is transmitted. Note that the LAST keyword is used to identify Part 2 as the final part to assemble:

```
XPRT 2 LAST


Please accept our offer for 10% off your next sprocket purchase.


.
250 Ok, mailmerge message enqueued
```

The "250 Ok, mailmerge message queued" notes that the message has been accepted.

Based on this example, recipient Jane User will receive this message:

```
From: Mr. Spacely <spacely@sprockets.com>

To: Jane User <jane@example.com>

Subject: Thanks for Being a Spacely Sprockets Customer


message date

Dear Jane,


Thank you for purchasing a red sprocket.


Please accept our offer for 10% off your next sprocket purchase.
```

Recipient Joe User will receive this message:

```
From: Mr. Spacely <spacely@sprockets.com>

To: Joe User <joe@example.com>

Subject: Thanks for Being a Spacely Sprockets Customer


message date

Dear Joe,


Thank you for purchasing a black sprocket.
```

Example Code

Cisco IronPort has created libraries in common programming languages to abstract the task of injecting IPMM messages into the Cisco IronPort appliance listener enabled for IPMM. Contact Cisco IronPort Customer Support for examples of how to use the IPMM library. The code is commented extensively to explain its syntax.

CHAPTER 17

The Cisco IronPort M-Series Security Management Appliance

The Cisco IronPort M-Series appliance is a special model of the Cisco IronPort appliances, specifically designed to serve as an external or “off box” spam quarantine for use with other Cisco IronPort appliances. This chapter discusses network planning, system setup, and general use of the Cisco IronPort M-Series appliance.

Overview

You can use an Cisco IronPort M-Series Security Management appliance to complement your Cisco IronPort Email Security appliance. The Cisco IronPort M-Series Security Management appliance is designed to serve as an external or “off box” location to monitor corporate policy settings and audit information. It combines hardware, an operating system (AsyncOS), and supporting services to centralize and consolidate important policy and runtime data, providing administrators and end users with a single interface for managing reporting and auditing information for the Cisco IronPort C-Series and X-Series Email Security appliances. The Cisco IronPort M-Series appliance ensures top performance from Cisco IronPort Email Security appliances, and protects corporate network integrity by increasing deployment flexibility. You can coordinate your security operations from a single Cisco IronPort M-Series appliance, or spread the load across multiple appliances.

The AsyncOS for Security Management appliance includes the following features:

- External Cisco IronPort Spam Quarantine. Hold spam and suspected spam messages for end users, and allow end users and administrators to review messages that are flagged as spam before making a final determination.
- Centralized reporting. Run reports on aggregated data from multiple Email Security appliances.
- Centralized tracking. Track email messages that traverse multiple Email Security appliances.

For information about configuring and using your Cisco IronPort Security Management appliance, see the *Cisco IronPort AsyncOS for Security Management User Guide* .

Network Planning

The Cisco IronPort M-Series appliance lets you separate the end user interfaces (mail applications, etc.) from the more secure gateway systems residing in your various DMZs. Using a two-layer firewall can provide you with flexibility in network planning so that end users will not connect directly to the outer DMZ (see Figure 17-1).

Figure 17-1 Typical Network Configuration Incorporating the Cisco IronPort M-Series Appliance

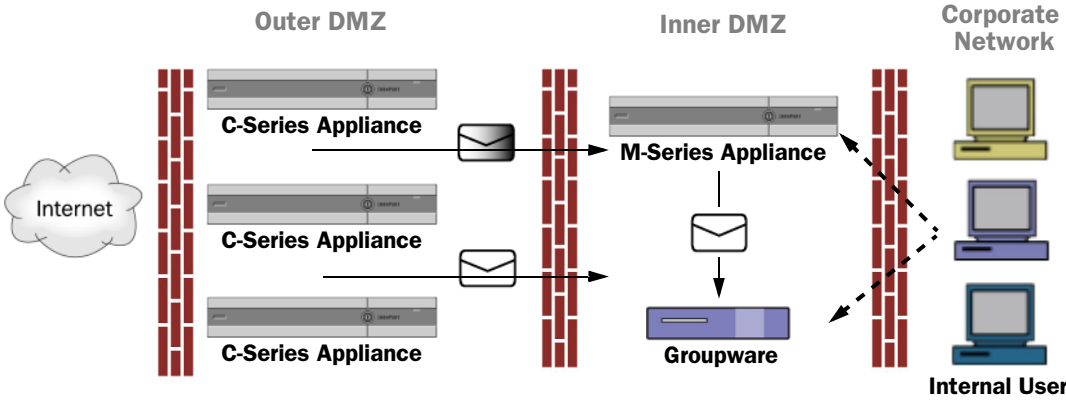

Large corporate data centers can share one Cisco IronPort M-Series appliance acting as an external Cisco IronPort Spam quarantine for one or more Cisco IronPort C- or X-Series appliances. Further, remote offices can be set up to maintain their own local Cisco IronPort appliance quarantines for local use (using the local Cisco IronPort Spam quarantine on C- or X-Series appliances).

Figure 17-1 shows a typical network configuration incorporating the Cisco IronPort M-Series appliance and multiple DMZs. Incoming mail from the Internet is received by the Cisco IronPort appliances in the outer DMZ. Clean mail is sent along to the MTA (groupware) in the inner DMZ and eventually to the end users within the corporate network.

Spam and suspected spam (depending on your mail flow policy settings) is sent to the Cisco IronPort M-Series appliance's Spam quarantine. End users may then access the quarantine and elect to delete spam and release messages they would like to have delivered to themselves. Messages remaining in the Cisco IronPort Spam quarantine are automatically deleted after a configurable amount of time (see the "Quarantines" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*).

Mail Flow and the Cisco IronPort M-Series Appliance

Mail is sent to the Cisco IronPort M-Series appliance from other Cisco IronPort (C- and X-Series) appliances. A Cisco IronPort appliance that is configured to send mail to a Cisco IronPort M-Series appliance will automatically expect to receive mail released from the M-Series appliance and will not re-process those messages when they are received back — messages will bypass the HAT and other policy or scanning settings and be delivered. For this to work, the IP address of the Cisco IronPort M-Series appliance must not change. If the IP address of the Cisco IronPort M-Series appliance changes, the receiving C- or X-Series appliance will process the message as it would any other incoming message. You should always use the same IP address for receiving and delivery on the Cisco IronPort M-Series appliance.

The Cisco IronPort M-Series appliance accepts mail for quarantining from the IP addresses specified in the Cisco IronPort Spam Quarantine settings. To configure the local quarantine on the Cisco IronPort M-Series appliance see the *Cisco IronPort AsyncOS for Security Management User Guide* . Note that the local quarantine on the Cisco IronPort M-Series appliance is referred to as an *external* quarantine by the other Cisco IronPort appliances sending mail to it.

Mail released by the Cisco IronPort M-Series appliance is delivered to the primary and secondary hosts (Cisco IronPort appliance or other groupware host) as defined in the Spam Quarantine Settings (see the *Cisco IronPort AsyncOS for Security Management User Guide*). Therefore, regardless of the number of Cisco IronPort appliances delivering mail to the Cisco IronPort M-Series appliance, all released mail, notifications, and alerts are sent to a single host (groupware or Cisco IronPort appliance). Take care to not overburden the primary host for delivery from the Cisco IronPort M-Series appliance.

Configuring Monitoring Services

Before you can use a Security Management appliance for centralized reporting and centralized tracking or as an external Cisco IronPort Spam Quarantine, you need to configure the monitoring services on the Email Security appliances.

When you configure the monitoring services on the Email Security appliances, you must also enable the services on the Security Management appliance. For more information, see the *Cisco IronPort AsyncOS for Security Management User Guide* .

You use monitoring services to run reports on email traffic, track message routing, and deliver suspect and spam messages to an external Cisco IronPort Spam Quarantine. You can configure one or more of the following services:

- **Centralized Reporting.** For more information, see Configuring an Email Security Appliance to Use Centralized Reporting, page 17-3.
- **Centralized Tracking.** For more information, see Configuring an Email Security Appliance to Use Centralized Tracking, page 17-4.
- **Cisco IronPort Spam Quarantine.** For more information, see Configuring an Email Security Appliance to Use an External Cisco IronPort Spam Quarantine, page 17-5.

Configuring an Email Security Appliance to Use Centralized Reporting

You can configure centralized reporting on an Email Security appliance at any time. Typically, you configure centralized reporting after you enable the feature on a Security Management appliance.

Note Before enabling centralized reporting, ensure that sufficient disk space is allocated to that service.

Step 1 Click Security Services > Reporting.

The Reporting Service Settings page is displayed.

Figure 17-2 The Reporting Service Settings Page

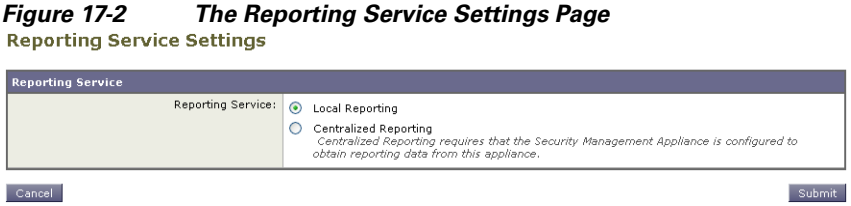

Step 2 In the Reporting Service section, select the Centralized Reporting option.

Step 3 Submit and commit your changes.

Note To use centralized reporting, you must enable the feature on the Email Security appliances *and* on the Security Management appliance. For information about enabling centralized reporting on the Security Management appliance, see the *Cisco IronPort AsyncOS for Security Management User Guide* .

Centralized Reporting Mode

After an Email Security appliance is configured to use centralized reporting and you add it to the Security Management appliance as a managed appliance, the Email Security appliance operates in centralized reporting mode. When an Email Security appliance is in centralized reporting mode, the scheduled reports for that appliance are suspended, and you can no longer access the scheduled report configuration page and the archived reports for the appliance. Also, the appliance stores only a week's worth of data. New data for the monthly and yearly reports is stored on the Security Management appliance. Existing data on the Email Security appliance for the monthly report is not transferred to the Security Management appliance. After centralized reporting is disabled, the Email Security appliance begins storing new monthly report data.

If you disable centralized reporting on the Email Security appliance, scheduled reports resume, and you can access its archived reports. After disabling centralized reporting, the appliance only displays data for the past hour and day, but not the past week or month. This is temporary. The appliance will display the reports for the past week and month after it accumulates enough data. If the Email Security appliance is placed back into centralized reporting mode, it will display data for the past week in the interactive reports.

Configuring an Email Security Appliance to Use Centralized Tracking

You can configure an Email Security appliance to use either local (on-box) tracking or centralized tracking.

Note You cannot enable both centralized and local tracking on an Email Security appliance.

Step 1 Click Security Services > Message Tracking.

The Message Tracking Service page is displayed.

Figure 17-3 The Message Tracking Service Page

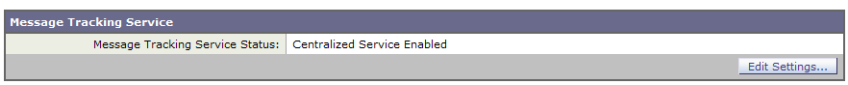

Step 2 In the Message Tracking Service section, click **Edit Settings**.

Figure 17-4 The Message Tracking Service Settings Page

Step 3 Select the Enable Message Tracking Service check box.

Step 4 Select the Centralized Tracking option.

Step 5 Optionally, select the check box to save information for rejected connections.

Note Saving tracking information for rejected connections can adversely affect the performance of the Security Management appliance.

Step 6 Submit and commit your changes.

Note To use centralized tracking, you must enable the feature on the Email Security appliances *and* the Security Management appliance. For information about enabling centralized tracking on the Security Management appliance, see the *Cisco IronPort AsyncOS for Security Management User Guide* .

Configuring an Email Security Appliance to Use an External Cisco IronPort Spam Quarantine

You need to enable the external spam quarantine feature on an Email Security appliance to use a Security Management appliance as an Cisco IronPort Spam Quarantine. You also need to provide the IP address and port number that the Email Security appliance uses to connect to the external spam quarantine.

Step 1 Click Security Services > External Spam Quarantine.

The External Spam Quarantine page is displayed.

Step 2 Click **Configure**.

The Configure External Spam Quarantine page is displayed.

Figure 17-5 The Configure External Spam Quarantine Page

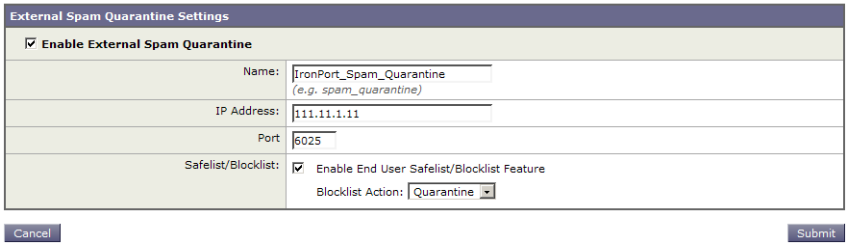

Step 3 In the External Spam Quarantine section, select the Enable External Spam Quarantine check box.

Step 4 In the Name field, enter the name of the Security Management appliance.

Step 5 Enter an IP address and port number. The IP address and port number for the Security Management appliance are configured on the Cisco IronPort Spam Quarantine page.

Step 6 Optionally, select the check box to enable the End User Safelist/Blocklist feature, and specify the appropriate blocklist action.

Step 7 Submit and commit your changes.

For more information about the Cisco IronPort Spam Quarantine and the End User Safelist/Blocklist feature, see "Quarantines" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide*. For more information about working with the Cisco IronPort Spam Quarantine on an M-Series appliance, see the *Cisco IronPort AsyncOS for Security Management User Guide* .

APPENDIX A

Accessing the Appliance

You can access any IP interface you create on the appliance through a variety of services.

By default, the following services are either enabled or disabled on each interface:

Table A-1 Services Enabled by Default on IP Interfaces

		Enabled by default?	
Service	**Default port**	**Management interface**	**New IP interfaces you create**
FTP	21	No	No
Telnet	23	Yes	No
SSH	22	Yes	No
HTTP	80	Yes	No
HTTPS	443	Yes	No

(a) the "Management Interface" settings shown here are also the default settings for the Data 1 Interface on Cisco IronPort C10/100 appliances.

- If you need to access the appliance via the graphical user interface (GUI), you must enable HTTP and/or HTTPS on an interface.
- If you need to access the appliance for the purposes of uploading or downloading configuration files, you must enable FTP or Telnet on an interface. See FTP Access, page A-4.
- You can also upload or download files using secure copy (`scp`).

IP Interfaces

An IP interface contains the network configuration data needed for an individual connection to the network. You can configure multiple IP interfaces to a physical Ethernet interface. You can also configure access to the Cisco IronPort Spam quarantine via an IP interface. For email delivery and Virtual Gateways, each IP interface acts as one Virtual Gateway address with a specific IP address and hostname. You can assign an Internet Protocol version 4 (IPv4) or version 6 (IPv6) to an IP inteface or both. You can also "join" interfaces into distinct groups (via the CLI), and the system will cycle through these groups when delivering email. Joining or grouping Virtual Gateways is useful for load-balancing large email campaigns across several interfaces. You can also create VLANs, and configure them just as you would any other interface (via the CLI). For more information, see the "Advanced Networking" chapter in the *Cisco IronPort AsyncOS for Email Advanced Configuration Guide*.

Figure A-1 IP Interfaces Page

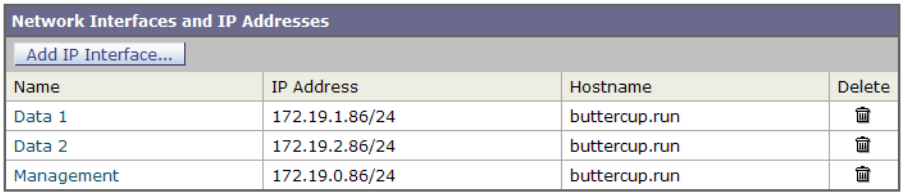

Configuring IP Interfaces

The Network > IP Interfaces page (and `interfaceconfig` command) allows you to add, edit, or delete IP interfaces.

Note You can not change the name or ethernet port associated with the Management interface on the M-Series appliance. Further, the Cisco IronPort M-Series appliance does not support all of the features discussed below (Virtual Gateways, for example).

The following information is required when you configure an IP interface:

Table A-2 IP Interface Components

Name	The nickname of the interface.
IPv4 address / Netmask	The IPv4 addresses within the same subnet cannot be configured on separate physical Ethernet interfaces. You can enter the subnetmask as a prefix in CIDR notation (e.g. `/24` for the 255.255.255.0 subnet).
IPv6 address / Prefix	The IPv6 addresses within the same subnet cannot be configured on separate physical Ethernet interfaces. IPv6 addresses must use leading zeros, such as `2001:0db8:85a3::8a2e:0370:7334`. You can enter the subnetmask is a prefix in CIDR notation (e.g. `2001:0db8:85a3::8a2e:0370:7334/64`.)
Broadcast address	Cisco IronPort AsyncOS automatically calculates the default broadcast address from the IP address and the netmask.
Hostname	The hostname that is related to the interface. This hostname will be used to identify the server during the SMTP conversation. You are responsible for entering a valid hostname associated with each IP address. The software does not check that DNS correctly resolves the hostname to the matching IP address, or that reverse DNS resolves to the given hostname.
Allowed services	FTP, SSH, Telnet, Cisco IronPort Spam Quarantine, HTTP, and HTTPS can be enabled or disabled on the interface. You can configure the port for each service. You can also specify the HTTP/HTTPS, port, and URL for the Cisco IronPort Spam Quarantine.

Note If you have completed the GUI's System Setup Wizard (or the Command Line Interface `systemsetup` command) as described in Chapter 3, "Setup and Installation" and committed the changes, one or two interfaces should already be configured on your appliance. (Refer to the settings you entered in the "Assign and Configure Logical IP Interface(s)" section.) In addition, the Management interface is configured on the Cisco IronPort appliance.

Creating IP Interfaces via the GUI

Step 1 Click **Add IP Interface** on the Network > IP Interfaces page. The Add IP Interface page is displayed:

Figure A-2 Add IP Interface Page

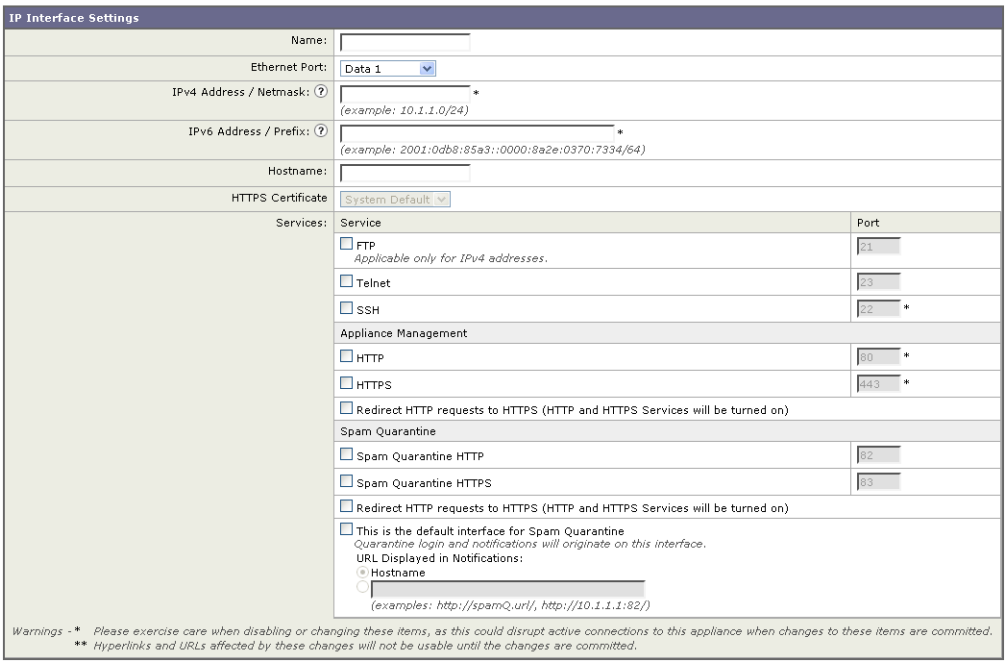

Step 2 Enter a name for the interface.

Step 3 Select an Ethernet port.

Step 4 Enter an IP address. You can enter an IPv4 address, an IPv6 address, or both. You can enter a subnetmask for the addresses using a CIDR format prefix, such as `/24` for the for IPv4 or `/64` for IPv6. If you enter both IPv4 and IPv6 addresses, the interface will use the version appropriate for each connection.

Step 5 Enter a hostname for the interface.

Step 6 Select a TLS certificate for HTTPS services.

Step 7 Mark the checkbox next to each service you wish to enable on this IP interface. Change the corresponding port if necessary.

Step 8 Select whether or not to enable redirecting HTTP to HTTPS for appliance management on the interface.

Step 9 If you are using the Cisco IronPort Spam quarantine, you can select HTTP or HTTPS or both and specify the port numbers for each. You can also select whether to redirect HTTP requests to HTTPS. Finally, you can specify whether the IP interface is the default interface for the Cisco IronPort Spam quarantine, whether to use the hostname as the URL, or provide a custom URL.

Step 10 Click **Submit**.

Step 11 Click the **Commit Changes** button, add an optional comment if necessary, and then click **Commit Changes** to finish creating the IP interface.

FTP Access

Warning **By disabling services via the Network > IP Interfaces page or the interfaceconfig command, you have the potential to disconnect yourself from the GUI or CLI, depending on how you are connected to the appliance. Do not disable services with this command if you are not able to reconnect to the appliance using another protocol, the Serial interface, or the default settings on the Management port.**

Step 1 Use the Network > IP Interfaces page (or the `interfaceconfig` command) to enable FTP access for the interface. The interface must have an IPv4 address in order to be accessed using FTP.

In this example, the Management interface is edited to enable FTP access on port 21 (the default port):

Figure A-3 Edit IP Interface Page

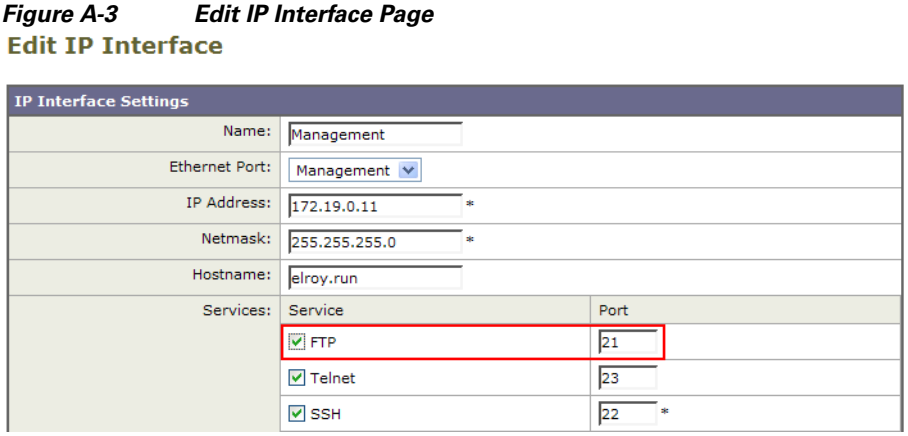

Note Remember to `commit` your changes before moving on to the next step.

Step 2 Access the interface via FTP. Ensure you are using the correct IP address for the interface. For example:
`ftp 192.168.42.42`

Many browsers also allow you to access interfaces via FTP. For example:
`ftp://192.10.10.10`

Note The FTP service on an appliance only uses IPv4 addresses, not IPv6 addresses.

Step 3 Browse to the directory for the specific task you are trying to accomplish. After you have accessed an interface via FTP, you can browse the following directories to copy and add ("GET" and "PUT") files. See Table A-2 on page A-5.

Table A-3 Directories available for access

Directory Name	Description
/antivirus	The directory where the Sophos Anti-Virus engine log files are kept. You can inspect the log files this directory to manually check for the last successful download of the virus definition file (`scan.dat`).
/avarchive **/bounces** **/cli_logs** **/delivery** **/error_logs** **/ftpd_logs** **/gui_logs** **/mail_logs** **/rptd_logs** **/sntpd.logs** **/status** **/system_logs**	Created automatically for **logging** via the System Administration > Logging page or the `logconfig` and `rollovernow` commands. See the "Logging" chapter in the *Cisco IronPort AsyncOS for Email Daily Management Guide* for a detailed description of each log. See "Log File Type Comparison" in the Logging chapter for the differences between each log file type.

Table A-3 Directories available for access (Continued)

Directory Name	Description
/MFM	The Mail Flow Monitoring database directory contains data for the Mail Flow Monitor functionality available from the GUI. Each subdirectory contains a README file that documents the record format for each file. You can copy these files to a different machine for record keeping, or load the files into a database and create your own analysis application. The record format is the same for all files in all directories; this format may change in future releases.
/saved_reports	The directory where all archived reports configured on the system are stored.
/configuration	The directory where data from the following pages and commands is exported to and/or imported (saved) from: • Virtual Gateway mappings (`altsrchost`) • configuration data in XML format (`saveconfig, loadconfig`) • Host Access Table (HAT) Page (`hostaccess`) • Recipient Access Table (RAT) Page (`rcptaccess`) • SMTP Routes Page (`smtproutes`) • alias tables (`aliasconfig`) • masquerading tables (`masquerade`) • message **filters** (`filters`) • global unsubscribe data (`unsubscribe`) • test messages for the `trace` command

Step 4 Use your FTP program to upload and download files to and from the appropriate directory.

Secure Copy (scp) Access

If your client operating system supports a secure copy (`scp`) command, you can copy files to and from the directories listed in Table A-2. For example, in the following example, the file `/tmp/test.txt` is copied from the client machine to the configuration directory of the appliance with the hostname of `mail3.example.com`.

Note that the command prompts for the password for the user (`admin`). This example is shown for reference only; your particular operating system's implementation of secure copy may vary.

```
% scp /tmp/test.txt admin@mail3.example.com:configuration

The authenticity of host 'mail3.example.com (192.168.42.42)' can't be established.

DSA key fingerprint is 69:02:01:1d:9b:eb:eb:80:0c:a1:f5:a6:61:da:c8:db.

Are you sure you want to continue connecting (yes/no)? yes

Warning: Permanently added 'mail3.example.com ' (DSA) to the list of known hosts.
```

```
admin@mail3.example.com's password: (type the password)

test.txt              100% |***************************|  1007       00:00

%
```

In this example, the same file is copied from the appliance to the client machine:

```
% scp admin@mail3.example.com:configuration/text.txt .

admin@mail3.example.com's password: (type the password)

test.txt              100% |***************************|  1007       00:00
```

You can use secure copy (scp) as an alternative to FTP to transfer files to and from the Cisco IronPort appliance.

Note Only users in the operators and administrators group can use secure copy (scp) to access the appliance. For more information, see information on adding users in "Common Admistrative Tasks" in *Cisco IronPort AsyncOS for Email Daily Management Guide*.

Accessing via a Serial Connection

If you are connecting to the appliance via a serial connection (see Connecting to the Appliance, page 3-9), Figure A-4 illustrates the pin numbers for the serial port connector, and Table A-4 defines the pin assignments and interface signals for the serial port connector.

Figure A-4 Pin Numbers for the Serial Port

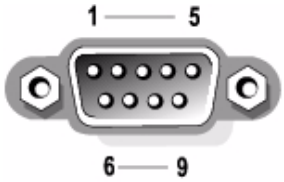

Table A-4 Serial Port Pin Assignments

Pin	Signal	I/O	Definition
1	DCD	I	Data carrier detect
2	SIN	I	Serial input
3	SOUT	O	Serial output
4	DTR	O	Data terminal ready
5	GND	n/a	Signal ground
6	DSR	I	Data set ready
7	RTS	I	Request to send

Table A-4 Serial Port Pin Assignments (Continued)

Pin	Signal	I/O	Definition
8	CTS	O	Clear to send
9	RI	I	Ring indicator
Shell	n/a	n/a	Chassis ground

APPENDIX B

Assigning Network and IP Addresses

This appendix describes general rules on networks and IP address assignments, and it presents some strategies for connecting the Cisco IronPort appliance to your network.

Ethernet Interfaces

The Cisco IronPort X1050/1060/1070, C650/660/670, and C350/360/370 appliances are equipped with as many as four Ethernet interfaces located on the rear panel of the system depending on the configuration (whether or not you have the optional optical network interface). They are labeled:

- Management
- Data1
- Data2
- Data3
- Data4

The Cisco IronPort C150/160 appliance is equipped with two Ethernet interfaces located on the rear panel of the system. They are labeled:

- Data1
- Data2

Selecting IP Addresses and Netmasks

When you configure the network, the Cisco IronPort appliance must be able to uniquely select an interface to send an outgoing packet. This requirement will drive some of the decisions regarding IP address and netmask selection for the Ethernet interfaces. The rule is that only one interface can be on a single network (as determined through the applications of netmasks to the IP addresses of the interfaces).

An IP address identifies a physical interface on any given network. A physical Ethernet interface can have more than one IP address for which it accepts packets. An Ethernet interface that has more than one IP address can send packets over that interface with any one of the IP addresses as the source address in the packet. This property is used in implementing Virtual Gateway technology.

The purpose of a netmask is to divide an IP address into a network address and a host address. The network address can be thought of as the network part (the bits matching the netmask) of the IP address. The host address is the remaining bits of the IP address. The number of bits in a four octet address that are significant are sometimes expressed in CIDR (Classless Inter-Domain Routing) style. This is a slash followed by the number of bits (1-32).

A netmask can be expressed in this way by simply counting the ones in binary, so `255.255.255.0` becomes "`/24`" and `255.255.240.0` becomes "`/20`".

Sample Interface Configurations

This section shows sample interface configurations based on some typical networks. The example will use two interfaces called Int1 and Int2. In the case of the Cisco IronPort appliance, these interface names can represent any two interfaces out of the three Cisco IronPort interfaces (Management, Data1, Data2).

Network 1:

Separate interfaces must appear to be on separate networks.

Interface	IP address	netmask	net address
Int1	`192.168.1.10`	`255.255.255.0`	`192.168.1.0/24`
Int2	`192.168.0.10`	`255.255.255.0`	`192.168.0.0/24`

Data addressed to `192.168.1.X` (where X is any number 1-255, except for your own address, 10 in this case) will go out on Int1. Anything addressed to `192.168.0.X` will go out on Int2. Any packet headed for some other address not in these formats, most likely out on a WAN or the Internet, will be sent to the default gateway which must itself be on one of these networks. The default gateway will then forward the packet on.

Network 2:

The network addresses (network parts of the IP addresses) of two different interfaces cannot be the same.

Ethernet Interface	IP address	netmask	net address
Int1	`192.168.1.10`	`255.255.0.0`	`192.168.0.0/16`
Int2	`192.168.0.10`	`255.255.0.0`	`192.168.0.0/16`

This situation presents a conflict in that two different Ethernet interfaces have the same network address. If a packet from the Cisco IronPort appliance is sent to `192.168.1.11`, there is no way to decide which Ethernet interface should be used to deliver the packet. If the two Ethernet interfaces are connected to two separate physical networks, the packet may be delivered to the incorrect network and never find its destination. The Cisco IronPort appliance will not allow you to configure your network with conflicts.

You can connect two Ethernet interfaces to the same physical network, but you must construct IP addresses and netmasks to allow the Cisco IronPort appliance to select a unique delivery interface.

IP Addresses, Interfaces, and Routing

When selecting an interface on which to perform a command or function in the GUI or CLI that allows you to select an interface (for example, upgrading AsyncOS, or configuring DNS, etc.), routing (your default gateway) will take precedence over your selection.

For example, suppose you have an Cisco IronPort appliance with the 3 network interfaces configured, each on a different network segment (assume all /24):

Ethernet	IP
Management	192.19.0.100
data1	192.19.1.100
data2	192.19.2.100

And your Default gateway is 192.19.0.1.

Now, if you perform an AsyncOS upgrade (or other command or function that allows you to select an interface) and you select the IP that is on data1 (192.19.1.100), you would expect all the TCP traffic to occur over the data1 ethernet interface. However, what happens is that the traffic will go out of the interface that is set as your default gateway, in this case Management, but will be stamped with the source address of the IP on data1.

Summary

The Cisco IronPort appliance must always be able to identify a unique interface over which a packet will be delivered. To make this decision, the Cisco IronPort appliance uses a combination of the packet's destination IP address, and the network and IP address settings of its Ethernet interfaces. The following table summarizes the preceding examples:

	Same Network	Different Network
Same Physical Interface	Allowed	Allowed
Different Physical Interface	Not Allowed	Allowed

Strategies for Connecting Your Cisco IronPort Appliance

Keep these things in mind when connecting your Cisco IronPort appliance:

- Administrative traffic (CLI, web interface, log delivery) traffic is usually small compared to email traffic.
- If two Ethernet interfaces are connected to the same network switch, but end up talking to a single interface on another host downstream, or are connected to a network hub where all data are echoed to all ports, no advantage is gained by using two interfaces.
- SMTP conversations over an interface operating at 1000Base-T will be slightly faster than ones over the same interfaces operating at 100Base-T, but only under ideal conditions.
- There is no point in optimizing connections to your network if there is a bottleneck in some other part of your delivery network. Bottlenecks most often occur in the connection to the Internet and further upstream at your connectivity provider.

The number of Cisco IronPort appliance interfaces that you choose to connect and how you address them should be dictated by the complexity of your underlying network. It is not necessary to connect multiple interfaces if your network topology or data volumes do not call for it. It is also possible to keep the connection simple at first as you familiarize yourself with the gateway and then increase the connectivity as volume and network topology require it.

APPENDIX C

Firewall Information

The following table lists the possible ports that may need to be opened for proper operation of the Cisco IronPort appliance (these are the default values).

Table C-1 Firewall Ports

Port	Protocol	In/Out	Hostname	Description
20/21	TCP	In or Out	AsyncOS IPs, FTP Server	FTP for aggregation of log files.
22	TCP	In	AsyncOS IPs	SSH access to the CLI, aggregation of log files.
22	TCP	Out	SSH Server	SSH aggregation of log files.
22	TCP	Out	SCP Server	SCP Push to log server
23	Telnet	In	AsyncOS IPs	Telnet access to the CLI, aggregation of log files.
23	Telnet	Out	Telnet Server	Telnet upgrades, aggregation of log files (not recommended).
25	TCP	Out	Any	SMTP to send email.
25	TCP	In	AsyncOS IPs	SMTP to receive bounced email or if injecting email from outside firewall.
80	HTTP	In	AsyncOS IPs	HTTP access to the GUI for system monitoring.
80	HTTP	Out	downloads.ironport.com	Service updates, except for AsyncOS upgrades and McAfee definitions.
80	HTTP	Out	updates.ironport.com	AsyncOS upgrades and McAfee Anti-Virus definitions.
80	HTTP	Out	cdn-microupdates.cloudmark.com	Used for updates to third-party spam component in Intelligent MultiScan. Appliance must also connect to CIDR range 208.83.136.0/22 for third-party phone home updates.
82	HTTP	In	AsyncOS IPs	Used for viewing the Cisco IronPort Anti-Spam quarantine.
83	HTTPS	In	AsyncOS IPs	Used for viewing the Cisco IronPort Anti-Spam quarantine.

Table C-1 Firewall Ports (Continued)

53	UDP/TCP	In & Out	DNS Servers	DNS if configured to use Internet root servers or other DNS servers outside the firewall. Also for SenderBase queries.
110	TCP	Out	POP Server	POP authentication for end users for Cisco IronPort Spam Quarantine
123	UDP	In & Out	NTP Server	NTP if time servers are outside firewall.
143	TCP	Out	IMAP Server	IMAP authentication for end users for Cisco IronPort Spam Quarantine
161	UDP	In	AsyncOS IPs	SNMP Queries
162	UDP	Out	Management Station	SNMP Traps
389 3268	LDAP	Out	LDAP Servers	LDAP if LDAP directory servers are outside firewall. LDAP authentication for Cisco IronPort Spam Quarantine
636 3269	LDAPS	Out	LDAPS	LDAPS — ActiveDirectory's Global Catalog Server
443	TCP	In	AsyncOS IPs	Secure HTTP (`https`) access to the GUI for system monitoring.
443	TCP	Out	res.cisco.com	Cisco Registered Envelope Service
443	TCP	Out	updates-static.ironport.com	Verify the latest files for the update server.
443	TCP	Out	phonehome.senderbase.org	Receive/Send Outbreak Filters
514	UDP/TCP	Out	Syslog Server	Syslog logging
628	TCP	In	AsyncOS IPs	QMQP if injecting email from outside firewall.
2222	CCS	In & Out	AsyncOS IPs	Cluster Communication Service (for Centralized Management).
6025	TCP	Out	AsyncOS IPs	Cisco IronPort Spam Quarantine

APPENDIX D

Cisco IronPort End User License Agreement

Cisco IronPort Systems, LLC Software License Agreement

NOTICE TO ALL USERS: CAREFULLY READ THE FOLLOWING LEGAL AGREEMENT ("AGREEMENT") FOR THE LICENSE OF THE SOFTWARE (AS DEFINED BELOW). BY CLICKING THE ACCEPT BUTTON OR ENTERING "Y" WHEN PROMPTED, YOU (EITHER AN INDIVIDUAL OR A SINGLE ENTITY, COLLECTIVELY, THE "COMPANY") CONSENT TO BE BOUND BY AND BECOME A PARTY TO THE FOLLOWING AGREEMENT BETWEEN CISCO IRONPORT SYSTEMS, LLC, A DELAWARE CORPORATION ("IRONPORT") AND COMPANY (COLLECTIVELY, THE "PARTIES"). BY CLICKING THE ACCEPT BUTTON OR ENTERING "Y" WHEN PROMPTED, YOU REPRESENT THAT (A) YOU ARE DULY AUTHORIZED TO REPRESENT YOUR COMPANY AND (B) YOU ACCEPT THE TERMS AND CONDITIONS OF THIS AGREEMENT ON BEHALF OF YOUR COMPANY, AND AS SUCH, AN AGREEMENT IS THEN FORMED. IF YOU OR THE COMPANY YOU REPRESENT (COLLECTIVELY, "COMPANY") DO NOT AGREE TO THE TERMS AND CONDITIONS OF THIS AGREEMENT, CLICK THE CANCEL BUTTON OR ENTER "N" WHEN PROMPTED AND PROMPTLY (BUT NO LATER THAT THIRTY (30) DAYS OF THE DELIVERY DATE, AS DEFINED BELOW) NOTIFY IRONPORT, OR THE RESELLER FROM WHOM YOU RECEIVED THE SOFTWARE, FOR A FULL REFUND OF THE PRICE PAID FOR THE SOFTWARE.

1. DEFINITIONS

1.1 "Company Service" means the Company's email or internet services provided to End Users for the purposes of conducting Company's internal business and which are enabled via Company's products as described in the purchase agreement, evaluation agreement, beta or pre-release agreement, purchase order, sales quote or other similar agreement between the Company and IronPort or its reseller ("Agreement") and the applicable user interface and IronPort's standard system guide documentation that outlines the system architecture and its interfaces (collectively, the "License Documentation").

1.2 "End User" means the employee, contractor or other agent authorized by Company to access to the Internet or use email services via the Company Service.

1.3 "Service(s)" means (i) the provision of the Software functionality, including Updates and Upgrades, and (ii) the provision of support by IronPort or its reseller, as the case may be.

1.4 “Software” means: (i) IronPort’s proprietary software licensed by IronPort to Company along with IronPort’s hardware products; (ii) any software provided by IronPort’s third-party licensors that is licensed to Company to be implemented for use with IronPort’s hardware products; (iii) any other IronPort software module(s) licensed by IronPort to Company along with IronPort’s hardware products; and (iv) any and all Updates and Upgrades thereto.

1.5 “Updates” means minor updates, error corrections and bug fixes that do not add significant new functions to the Software, and that are released by IronPort or its third party licensors. Updates are designated by an increase to the Software’s release number to the right of the decimal point (e.g., Software 1.0 to Software 1.1). The term Updates specifically excludes Upgrades or new software versions marketed and licensed by IronPort or its third party licensors as a separate product.

1.6 “Upgrade(s)” means revisions to the Software, which add new enhancements to existing functionality, if and when it is released by IronPort or its third party licensors, in their sole discretion. Upgrades are designated by an increase in the Software’s release number, located to the left of the decimal point (e.g., Software 1.x to Software 2.0). In no event shall Upgrades include any new versions of the Software marketed and licensed by IronPort or its third party licensors as a separate product.

2. LICENSE GRANTS AND CONSENT TO TERMS OF DATA COLLECTION

2.1 License of Software. By using the Software and the License Documentation, Company agrees to be bound by the terms of this Agreement, and so long as Company is in compliance with this Agreement, IronPort hereby grants to Company a non-exclusive, non-sublicensable, non-transferable, worldwide license during the Term to use the Software only on IronPort’s hardware products, solely in connection with the provision of the Company Service to End Users. The duration and scope of this license(s) is further defined in the License Documentation. Except as expressly provided herein, no right, title or interest in any Software is granted to the Company by IronPort, IronPort’s resellers or their respective licensors. This license and any Services are co-terminus.

2.2 Consent and License to Use Data. Subject to Section 8 hereof, and subject to the IronPort Privacy Statement at http://www.IronPort.com/privacy.html, as the same may be amended from time to time by IronPort with notice to Company, Company hereby consents and grants to IronPort a license to collect and use the data from the Company as described in the License Documentation, as the same may be updated from time to time by IronPort (“Data”). To the extent that reports or statistics are generated using the Data, they shall be disclosed only in the aggregate and no End User identifying information may be surmised from the Data, including without limitation, user names, phone numbers, unobfuscated file names, email addresses, physical addresses and file content. Notwithstanding the foregoing, Company may terminate IronPort’s right to collect and use Data at any time upon prior written or electronic notification, provided that the Software or components of the Software may not be available to Company if such right is terminated.

3. CONFIDENTIALITY. Each Party agrees to hold in confidence all Confidential Information of the other Party to the same extent that it protects its own similar Confidential Information (and in no event using less than a reasonable degree of care) and to use such Confidential Information only as permitted under this Agreement. For purposes of this Agreement “Confidential Information” means information of a party marked “Confidential” or information reasonably considered by the disclosing Party to be of a proprietary or confidential nature; provided that the Data, the Software, information disclosed in design reviews and any pre-production releases of the Software provided by IronPort is expressly designated Confidential Information whether or not marked as such.

4. PROPRIETARY RIGHTS; OWNERSHIP. Title to and ownership of the Software and other materials and all associated Intellectual Property Rights (as defined below) related to the foregoing provided by IronPort or its reseller to Company will remain the exclusive property of IronPort and/or its superior licensors. Company and its employees and agents will not remove or alter any trademarks, or other proprietary notices, legends, symbols, or labels appearing on or in copies of the Software or other materials delivered to Company by IronPort or its reseller. Company will not modify, transfer, resell for

profit, distribute, copy, enhance, adapt, translate, decompile, reverse engineer, disassemble, or otherwise determine, or attempt to derive source code for any Software or any internal data files generated by the Software or to create any derivative works based on the Software or the License Documentation, and agrees not to permit or authorize anyone else to do so. Unless otherwise agreed in writing, any programs, inventions, concepts, documentation, specifications or other written or graphical materials and media created or developed by IronPort or its superior licensors during the course of its performance of this Agreement, or any related consulting or professional service agreements, including all copyrights, database rights, patents, trade secrets, trademark, moral rights, or other intellectual property rights ("Intellectual Property Right(s)") associated with the performance of such work shall belong exclusively to IronPort or its superior licensors and shall, in no way be considered a work made for hire for Company within the meaning of Title 17 of the United States Code (Copyright Act of 1976).

5. LIMITED WARRANTY AND WARRANTY DISCLAIMERS

5.1 Limited Warranty. IronPort warrants to Company that the Software, when properly installed and properly used, will substantially conform to the specifications in the License Documentation for a period of ninety (90) days from the delivery date or the period set forth in the License Documentation, whichever is longer ("Warranty Period"). FOR ANY BREACH OF THE WARRANTY CONTAINED IN THIS SECTION, COMPANY'S EXCLUSIVE REMEDY AND IRONPORT'S ENTIRE LIABILITY, WILL BE PROMPT CORRECTION OF ANY ERROR OR NONCONFORMITY, PROVIDED THAT THE NONCONFORMITY HAS BEEN REPORTED TO IRONPORT AND/OR ITS RESELLER BY COMPANY WITHIN THE WARRANTY PERIOD. THIS WARRANTY IS MADE SOLELY TO COMPANY AND IS NOT TRANSFERABLE TO ANY END USER OR OTHER THIRD PARTY. IronPort shall have no liability for breach of warranty under this Section or otherwise for breach of this Agreement if such breach arises directly or indirectly out of or in connection with the following: (i) any unauthorized, improper, incomplete or inadequate maintenance or calibration of the Software by Company or any third party; (ii) any third party hardware software, services or system(s); (iii) any unauthorized modification or alteration of the Software or Services; (iv) any unauthorized or improper use or operation of the Software or Company's failure to comply with any applicable environmental specification; or (v) a failure to install and/or use Updates, Upgrades, fixes or revisions provided by IronPort or its resellers from time to time.

5.2 WARRANTY DISCLAIMER. THE EXPRESS WARRANTIES SET FORTH IN SECTION 5.1 OF THIS AGREEMENT CONSTITUTE THE ONLY PERFORMANCE WARRANTIES WITH RESPECT TO THE SOFTWARE OR SERVICES. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IRONPORT LICENSES THE SOFTWARE AND SERVICES HEREUNDER ON AN "AS IS" BASIS. EXCEPT AS SPECIFICALLY SET FORTH HEREIN, IRONPORT AND ITS SUPERIOR LICENSORS MAKE NO REPRESENTATIONS OR WARRANTIES OF ANY KIND, WHETHER EXPRESS, IMPLIED, OR STATUTORY (EITHER IN FACT OR BY OPERATION OF LAW), AND EXPRESSLY DISCLAIM ALL OTHER WARRANTIES, INCLUDING WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. NEITHER IRONPORT NOR ITS THIRD PARTY LICENSORS WARRANT THAT THE SOFTWARE OR SERVICES (1) IS FREE FROM DEFECTS, ERRORS OR BUGS, (2) THAT OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED, OR (3) THAT ANY RESULTS OR INFORMATION THAT IS OR MAY BE DERIVED FROM THE USE OF THE SOFTWARE WILL BE ACCURATE, COMPLETE, RELIABLE AND/OR SECURE.

6. LIMITATION OF LIABILITY. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IN NO EVENT WILL EITHER PARTY BE LIABLE TO THE OTHER FOR ANY LOSS OF PROFITS, COSTS OF PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, LOSS OF BUSINESS, LOSS OF USE OR DATA, INTERRUPTION OF BUSINESS, OR FOR INDIRECT, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES OF ANY KIND, EVEN IF SUCH PARTY RECEIVED ADVANCE NOTICE OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL THE LIABILITY OF EITHER PARTY ARISING UNDER ANY PROVISION OF THIS AGREEMENT, REGARDLESS OF WHETHER THE CLAIM FOR SUCH DAMAGES IS

BASED IN CONTRACT, TORT, OR OTHER LEGAL THEORY, EXCEED THE TOTAL AMOUNT PAID FOR THE SOFTWARE OR SERVICES DURING THE TWELVE (12) MONTHS PRIOR TO THE EVENT GIVING RISE TO SUCH LIABILITY.

7. TERM AND TERMINATION. The term of this Agreement shall be as set forth in the License Documentation (the "Term"). If IronPort defaults in the performance of any material provision of this Agreement or the License Documentation, then Company may terminate this Agreement upon thirty (30) days written notice if the default is not cured during such thirty (30) day period. If Company defaults in the performance of any material provision of this Agreement or the License Documentation, IronPort may terminate this Agreement upon thirty (30) days written notice if the default is not cured during such thirty (30) day notice and without a refund. This Agreement may be terminated by one Party immediately at any time, without notice, upon (i) the institution by or against the other Party of insolvency, receivership or bankruptcy proceedings or any other proceedings for the settlement of such Party's debts, (ii) such other Party making a general assignment for the benefit of creditors, or (iii) such other Party's dissolution. The license granted in Section 2 will immediately terminate upon this Agreement's termination or expiration. Within thirty (30) calendar days after termination or expiration of this Agreement, Company will deliver to IronPort or its reseller or destroy all copies of the Software and any other materials or documentation provided to Company by IronPort or its reseller under this Agreement.

8. U.S. GOVERNMENT RESTRICTED RIGHTS; EXPORT CONTROL. The Software and accompanying License Documentation are deemed to be "commercial computer software" and "commercial computer software documentation," respectively, pursuant to DFAR Section 227.7202 and FAR Section 12.212, as applicable. Any use, modification, reproduction, release, performance, display or disclosure of the Software and accompanying License Documentation by the United States Government shall be governed solely by the terms of this Agreement and shall be prohibited except to the extent expressly permitted by the terms of this Agreement. Company acknowledges that the Software and License Documentation must be exported in accordance with U.S. Export Administration Regulations and diversion contrary to U.S. laws is prohibited. Company represents that neither the United States Bureau of Export Administration nor any other federal agency has suspended, revoked or denied Company export privileges. Company represents that Company will not use or transfer the Software for end use relating to any nuclear, chemical or biological weapons, or missile technology unless authorized by the U.S. Government by regulation or specific license. Company acknowledges it is Company's ultimate responsibility to comply with any and all import and export restrictions, and other applicable laws, in the U.S. or elsewhere, and that IronPort or its reseller has no further responsibility after the initial sale to Company within the original country of sale.

9. MISCELLANEOUS. This Agreement is governed by the laws of the United States and the State of California, without reference to conflict of laws principles. The application of the United Nations Convention of Contracts for the International Sale of Goods is expressly excluded. Nothing contained herein shall be construed as creating any agency, partnership, or other form of joint enterprise between the parties. Neither party shall be liable hereunder by reason of any failure or delay in the performance of its obligations hereunder (except for the payment of money) on account of (i) any provision of any present or future law or regulation of the United States or any applicable law that applies to the subject hereof, and (ii) interruptions in the electrical supply, failure of the Internet, strikes, shortages, riots, insurrection, fires, flood, storm, explosions, acts of God, war, terrorism, governmental action, labor conditions, earthquakes, or any other cause which is beyond the reasonable control of such party. This Agreement and the License Documentation set forth all rights for the user of the Software and is the entire agreement between the parties and supersedes any other communications with respect to the Software and License Documentation. The terms and conditions of this Agreement will prevail, notwithstanding any variance with the License Documentation or any purchase order or other written instrument submitted by a party, whether formally rejected by the other party or not. This Agreement may not be modified except by a written addendum issued by a duly authorized representative of IronPort, except that IronPort may modify the IronPort Privacy Statement at any time, in its discretion, via notification to Company of such modification that will be posted at http://www.IronPort.com/privacy.html. No provision hereof shall be deemed waived unless such waiver

shall be in writing and signed by IronPort or a duly authorized representative of IronPort. If any provision of this Agreement is held invalid, the remainder of this Agreement shall continue in full force and effect. The parties confirm that it is their wish that this Agreement has been written in the English language only.

10. IRONPORT CONTACT INFORMATION. If Company wants to contact IronPort for any reason, please write to IronPort Systems, Inc., 950 Elm Avenue, San Bruno, California 94066, or call or fax us at tel: 650.989.6500 and fax: 650.989.6543.

GLOSSARY

A

Allowed Hosts — Computers that are allowed to relay email through the Cisco IronPort appliance via a private listener. Allowed hosts are defined by their hostnames or IP addresses.

Anti-Virus — Sophos and McAfee Anti-Virus scanning engines provide cross-platform anti-virus protection, detection and disinfection. through virus detection engines which scans files for viruses, Trojan horses and worms. These programs come under the generic term of *malware*, meaning "malicious software." The similarities between all types of malware allow anti-virus scanners to detect and remove not only viruses, but also all types of malicious software.

B

Blacklist — A list of known bad senders. By default, senders in the Blacklist sender group of a public listener are rejected by the parameters set in the $BLOCKED *mail flow policy.*

C

Character Set (Double-byte) — Double Byte Character Sets are foreign-language character sets requiring more than one byte of information to express each character.

CIDR Notation — Classless Inter-Domain Routing. A convenient shorthand for describing a range of IP addresses within their network contexts using an arbitrary number of bits. Using this notation, you note the network prefix part of an address by adding a forward slash (/) followed by the number of bits used for the network part. Thus a Class C network can be described in prefix notation as 192.168.0.1/24. A CIDR specification of 206.13.1.48/25 would include any address in which the first 25 bits of the address matched the first 25 bits of 206.13.1.48.

Content Filters — Content-based filters used to process messages during the Per-Recipient Scanning phase of the work queue in the email pipeline. Content filters are evoked after Message filters, and act on individual splintered messages.

Content Matching Classifier The detection component of the RSA data loss prevention scanning engine. A classifier contains a number of rules for detecting sensitive data, along with context rules that search for supporting data. For example, a credit card classifier not only requires that the message contain a string that matches a credit card number, but that it also contains supporting information such as an expiration data, a credit card company name, or an address.

Conversational Bounce A bounce that occurs within the SMTP conversation. The two types of conversational bounces are *hard bounces* and `soft bounces`.

D

Debounce Timeout The amount of time, in seconds, the system will refrain from sending the identical alert to the user.

Delayed Bounce A bounce that occurs within the SMTP conversation. The recipient host accepts the message for delivery, only to bounce it at a later time.

Delivery The act of delivering email messages to recipient domains or internal mail hosts from the Cisco IronPort appliance from a specific IP interface. The Cisco IronPort appliance can deliver messages from multiple IP interfaces within same physical machine using Virtual Gateway technology. Each Virtual Gateway contains a distinct IP address, hostname and domain, and email queue, and you can configure different mail flow policies and scanning strategies for each.

You can tailor the configuration of the delivery that the Cisco IronPort appliance performs, including the maximum simultaneous connections to remote hosts, the per-Virtual Gateway limit of maximum simultaneous connections to the host, and whether the conversations to remote hosts are encrypted.

DLP Data loss prevention. RSA Security DLP scanning engine protects your organization's information and intellectual property and enforces regulatory and organizational compliance by preventing users from unintentionally emailing sensitive data.

DLP Incident A data loss prevention incident occurs when a DLP policy detects one or more DLP violations that merit attention in an outgoing message.

DLP Policy A data loss prevention policy is a set of conditions used to determine whether an outgoing message contains sensitive data and the actions that AsyncOS takes on a message that contains such data.

DLP Risk Factor A score of 0 to 100 that represents the security risk of the DLP violations detected in an outgoing message. Based on the risk factor, the DLP policy determines the actions to take on the message.

DLP Violation An instance of data being found in a message that violates your organization's DLP rules.

DNS Domain Name System. See RFC 1045 and RFC 1035. DNS servers on a network resolve IP addresses to hostnames, and vice versa.

DoS attack Denial of Service attack, can also be in the form of DDos (Distributed Denial of Service Attack). An attack on a network or computer, the primary aim of which is to disrupt access to a given service.

DSN Delivery Status Notification, a bounced message.

E

Email Security Manager A single, comprehensive dashboard to manage all email security services and applications on IronPort appliances. Email Security Manager allows you to manage Outbreak Filters, Anti-Spam, Anti-Virus, and email content policies — on a per-recipient or per-sender basis, through distinct inbound and outbound policies. See also *Content Filters*.

Envelope Recipient The recipient of an email message, as defined in the RCPT TO: SMTP command. Also sometimes referred to as the "Recipient To" or "Envelope To" address.

Envelope Sender The sender of an email message, as defined in the MAIL FROM: SMTP command. Also sometimes referred to as the "Mail From" or "Envelope From" address.

F

False Negative A spam message or a message containing a virus or a DLP violation that was not detected as such.

False Positive A message falsely categorized as spam or as containing a virus or DLP violation.

Fully-Qualified Domain Name (FQDN) A domain name including all higher level domain names up to the top-level domain name; for example: `mail3.example.com` is a fully qualified domain name for the *host* at 192.168.42.42; `example.com` is the fully qualified domain name for the example.com *domain*. The fully qualified domain name must be unique within the Internet.

H

Hard Bounced Message A message that is permanently undeliverable. This can happen during the SMTP conversation or afterward.

HAT Host Access Table. The HAT maintains a set of rules that control incoming connections from remote hosts for a listener. Every *listener* has its own HAT. HATs are defined for public and private listeners, and contain *mail flow policies* and *sender groups*.

I

IDE File Virus Definition File. An IDE file contains signatures or definitions used by anti-virus software to detect viruses.

L

LDAP Lightweight Directory Access Protocol. A protocol used to access information about people (including email addresses), organizations, and other resources in an Internet directory or intranet directory.

Listener A listener describes an email processing service that will be configured on a particular IP interface. Listeners only apply to email entering the Cisco IronPort appliance — either from the internal systems within your network or from the Internet. IronPort 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 injector" or even a "SMTP daemon" running for each IP address you specify.

IronPort AsyncOS differentiates between *public* listeners — which by default have the characteristics for receiving email from the Internet — and *private* listeners that are intended to accept email only from internal (groupware, POP/IMAP, and other message generation) systems.

Log Subscription Creation of log files that monitor the performance of the Cisco IronPort appliance. The log files are stored in local disk(s) and can also be transferred and stored in a remote system. Typical attributes of a log subscription include: name, component to monitor (email operations, server), format, and transfer method.

M

Mail Flow Policies A mail flow policy is a way of expressing a group of *Host Access Table* (HAT) parameters (an access rule, followed by *rate limiting* parameters and custom SMTP codes and responses) for a *listener.* Together, *sender groups* and mail flow policies are defined in a listener's HAT. Your Cisco IronPort appliance ships with the predefined mail flow policies and sender groups for listeners.

MAIL FROM See *Envelope Sender.*

Maximum Number of Retries The maximum number of times that redelivery of a *soft bounced* message will be attempted before being *hard bounced*.

Maximum Time in Queue The maximum length of time that a *soft bounced* message will stay in the email queue for *delivery* before being *hard bounced*.

MTA Mail Transfer Agent, or Messaging Transfer Agent. The program responsible for accepting, routing, and delivering email messages. Upon receiving a message from a Mail User Agent or another MTA, the MTA stores a message temporarily locally, analyses the recipients, and routes it to another MTA (routing). It may edit and/or add to the message headers. The Cisco IronPort appliance is an MTA that combines hardware, a hardened operating system, application, and supporting services to produce a purpose-built, rack-mount server appliance dedicated for enterprise messaging.

MUA Mail User Agent. The program that allows the user to compose and read email messages. The MUA provides the interface between the user and the Message Transfer Agent. Outgoing mail is eventually handed over to an MTA for delivery.

MX Record Specifies the MTA on the Internet responsible for accepting mail for a specified domain. A Mail Exchange record creates a mail route for a domain name. A domain name can have multiple mail routes, each assigned a priority number. The mail route with the lowest number identifies the primary server responsible for the domain. Other mail servers listed will be used as backup.

N

Non-Conversational Bounce A bounce that occurs due to a message being returned after the message was accepted for delivery by the recipient host. These can be soft (4XX) or hard (5XX) bounces. You can analyze these bounce responses to determine what to do with the recipient messages (e.g. re-send soft bounced recipient messages and remove hard bounced recipients from database).

NTP Network Time Protocol. The `ntpconfig` command configures IronPort AsyncOS to use Network Time Protocol (NTP) to synchronize the system clock with other computers.

O

Open Relay An open relay (sometimes called an “insecure relay” or a “third party” relay) is an SMTP email server that allows unchecked third-party relay of email messages. By processing email that is neither for nor from a local user, an open relay makes it possible for an unknown senders to route large volumes of email (typically *spam*) through your *gateway*. The `listenerconfig` and `systemsetup` commands prevent you from unintentionally configuring your system as an open relay.

Outbreak Filters IronPort’s Outbreak Filters feature provides an additional layer of protection from viruses. The Outbreak Filters feature quarantines suspicious email messages, holding the messages until an updated virus IDE is available. or until they are deemed not a threat.

Q

Queue In the Cisco IronPort appliance, you can delete, bounce, suspend, or redirect messages in the email queue. This email queue of messages for destination domains is also referred to as the *delivery queue*. The queue of messages waiting to be processed by IronPort Anti-Spam or message filter actions is referred to as the *work queue*. You can view the status of both queues using the `status detail` command.

R

RAT Recipient Access Table. The Recipient Access Table defines which recipients will be accepted by a public listener. The table specifies the address (which may be a partial address or hostname) and whether to accept or reject it. You can optionally include the SMTP response to the RCPT TO command for that recipient. The RAT typically contains your local domains.

Rate Limiting Rate limiting limits the maximum number of messages per session, the maximum number of recipients per message, the maximum message size, the maximum recipients per hour, and the maximum number of concurrent connections you are willing to accept from a remote host.

RCPT TO See *Envelope Recipient*.

Receiving The act of receiving email messages on a specific listener configured on an IP interface. The Cisco IronPort appliance configures listeners to receive email messages — either inbound from the Internet, or outbound from your internal systems.

Reputation Filter A way of filtering suspicious senders based on their reputation. The SenderBase Reputation Service provides an accurate, flexible way for you to reject or "throttle" suspected *spam* based on the connecting IP address of the remote host.

S

Sender Group A sender group is simply a list of senders gathered together for the purposes of handling email from those senders in the same way (that is, applying a mail flow policy to a group of senders). A sender group is a list of senders (identified by IP address, IP range, host/domain, SenderBase Reputation Service classification, SenderBase Reputation score range, or DNS List query response) separated by commas in a listener's Host Access Table (HAT). You assign a name for sender groups, as well as *mail flow policies*.

Soft Bounced Message A message whose delivery will be reattempted at a later time base on the configured *maximum number of retries* or *maximum time in queue*.

Spam Unwanted, Unsolicited Commercial bulk Email (UCE/UBE). Anti-spam scanning identifies email messages that are suspected to be spam, according to its filtering rules.

STARTTLS Transport Layer Security (TLS) is an improved version of the Secure Socket Layer (SSL) technology. It is a widely used mechanism for encrypting SMTP conversations over the Internet. The IronPort AsyncOS operating system supports the STARTTLS extension to SMTP (Secure SMTP over TLS), described in RFC 2487.

T

TOC Threat Operations Center. This refers to all the staff, tools, data and facilities involved in detecting and responding to virus outbreaks.

W

Whitelist A list of known good senders. Add senders you trust to the Whitelist *sender group*. The $TRUSTED mail flow policy is configured so that email from senders you trust has no *rate limiting* enabled, and the content from those senders is not subject to anti-spam scanning.

INDEX

Symbols

Numerics

A

B

C

D

E

F

G

H

I

L

M

N

O

P

Q

R

S

T

U

V

W

X