



Cisco 7.0 Spam & Virus Blocker User Guide

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If the equipment causes interference to radio or television reception, which can be determined by turning the equipment off and on, users are encouraged to try to correct the interference by using one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

Modifications to this product not authorized by Cisco could void the FCC approval and negate your authority to operate the product.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The Cisco logo, IronPort Systems, Cisco v, Virus Outbreak Filters, Context Adaptive Scanning Engine (CASE), and SenderBase are trademarks of Cisco IronPort Systems LLC.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

Cisco 7.0 Spam & Virus Blocker User Guide
© 2010 Cisco Systems, Inc. All rights reserved.



CONTENTS

Preface ix

What's New in This Release ix

How to Use This Guide xi

How This Book is Organized xi

CHAPTER 1

Blocker Installation and Administration 1-1

Blocker Installation and Administration: Overview 1-1

Chapter Concepts/Glossary 1-1

Before You Begin 1-2

Blocker Pre-Installation Steps 1-2

Configuration Options 1-3

Register the Blocker Appliance in DNS 1-3

Network Configuration 1-3

Typical Configuration 1-4

DNS Settings 1-5

NAT Settings 1-6

Firewall Settings 1-7

Running the System Setup Wizard 1-9

Configuring Your System to Use LDAP Authentication 1-15

The Blocker Interface 1-16

Logging In 1-16

Saving Your Changes 1-17

The Configuration Summary Page 1-18

System Administration 1-21

Shutting Down the Blocker 1-21

Rebooting the Blocker 1-21

Open a Support Case 1-22

Upgrading the Blocker 1-22

Working with Feature Keys 1-23

The Feature Keys Page 1-23

Feature Key Settings 1-23

Editing Update Settings 1-23

Adding Users 1-24

Changing Your Password 1-25

- Configuring the Return Address for Various Generated Messages 1-25
- Alerts 1-26
- Configuring Domain Name System (DNS) Settings 1-26
- Configuring TCP/IP Traffic Routes 1-26
- System Time 1-26
- Managing the Configuration File 1-26
- Troubleshooting Tips 1-27

CHAPTER 2

Accepting Mail 2-29

- Accepting Mail: Overview 2-29
 - Chapter Concepts/Glossary 2-29
 - Before You Begin 2-30
- Understanding How the Blocker Accepts Mail (Public Listeners) 2-30
 - IP Interfaces 2-31
 - What if My Blocker Isn't at the Edge of My Network? 2-31
- Filtering Your Mail Based on the Sender 2-32
 - The Host Access Table (HAT) 2-32
 - What is an Open Relay and Why Should I Prevent It? 2-33
 - Mail Flow Policies: Access Rules and Parameters 2-33
- Understanding the Default Sender Groups and Mail Flow Policies 2-37
 - Understanding SBRS (SenderBase Reputation Service) 2-37
 - Implementing SenderBase Reputation Filters 2-39
 - Understanding the Email Pipeline 2-40
- Troubleshooting Tips 2-42

CHAPTER 3

Mail Routing For Users and Groups 3-43

- Creating Mail Rules for Groups of Users: Overview 3-43
- Chapter Concepts/Glossary 3-43
 - Before You Begin 3-44
- How Do I Create a List of Users? 3-44
- Defining the Users You Accept Mail For (Using RAT) 3-45
 - Adding New Mail Users or Groups of Users (in the RAT) 3-45
- Using LDAP to Route Mail for Users and Groups of Users 3-46
 - Configuring the Blocker to work with LDAP 3-47
- Understanding Mail Policies 3-48
 - Incoming vs. Outgoing Messages 3-49
 - Policy Matching 3-49
 - First Match Wins 3-50

Examples of Policy Matching	3-50
Creating a New Policy	3-51

CHAPTER 4**Sending Mail 4-53**

Sending Mail: Overview	4-53
Chapter Concepts/Glossary	4-53
Before You Begin	4-54
How Do I Send Mail From My Blocker?	4-54
Configuring the Listener to Send Mail From Your Network	4-54
Creating a Mail Flow Policy for Relaying Mail	4-55
Adding Disclaimers to Outgoing Mail	4-57
Disclaimer Text	4-57
Adding Disclaimer Text via a Listener	4-58
Using Email Tagging to Prevent Bounce Attacks	4-58
Configuring Bounce Verification	4-58

CHAPTER 5**Configuring Spam and Virus Protection 5-61**

Virus and Spam Protection: Overview	5-61
Chapter Concepts/Glossary	5-61
Enabling Virus Protection	5-62
Steps to Enable Virus Protection	5-62
Editing the Anti-Virus Settings for a Mail Policy	5-62
Anti-Virus Notification Templates	5-63
Working with System Quarantines	5-64
System Quarantine Settings	5-64
Retention Time	5-64
Default Action	5-65
Overflow Messages	5-65
Subject Tagging	5-65
Add X-Header	5-66
Strip Attachment	5-66
Users and User Groups	5-66
Creating System Quarantines	5-66
Editing System Quarantines	5-67
Deleting System Quarantines	5-67
Enabling Spam Protection	5-67
Steps to Enable Anti-Spam Scanning	5-68
Configuring the Spam Quarantine	5-68

- Enabling and Disabling the Spam Quarantine 5-69
- Disabling the Spam Quarantine 5-69
- Configuring Settings for the Spam Quarantine 5-69
- Enabling End User Safelists and Blocklists 5-70
- Creating and Maintaining Safelists and Blocklists 5-70
 - Administrator Tasks for Creating and Maintaining Safelists and Blocklists 5-71
- Enabling and Configuring Safelist/Blocklist Settings 5-71
 - Backing Up and Restoring the Safelist/Blocklist Database 5-71
 - Troubleshooting Safelists and Blocklists 5-72
 - End User Tasks for Configuring Safelists and Blocklists 5-72
- Adding Entries to Safelists 5-73
- Adding Entries to Blocklists 5-74

CHAPTER 6

Creating Custom Rules Using Content Filters 6-75

- Creating Custom Rules Using Content Filters: Overview 6-75
 - Chapter Concepts/Glossary 6-75
- Before You Begin 6-76
- Understanding How Content Filters Work 6-76
- How Do I Create a Content Filter? 6-77
 - How Do I Add Conditions to My Filter? 6-77
 - What Kinds of Actions Can I Configure? 6-82
 - Action Variables 6-86
- Understanding Threshold Scoring in Content Filtering 6-88
 - Threshold Scoring for Message Bodies and Attachments 6-89
 - Threshold Scoring Multipart/Alternative MIME Parts 6-89
- How Do I Create a List of Terms to Search For? 6-89
 - Adding Dictionaries 6-90
 - Matching Case-Sensitive Words 6-90
 - Matching Whole Words Only 6-90
 - Sorting Terms 6-91
 - Dictionary Entry Syntax 6-91
 - Deleting Dictionaries 6-91
 - Importing and Exporting Dictionaries 6-91

CHAPTER 7

Monitoring Your Blocker 7-93

- Monitoring Your Blocker: Overview 7-93
 - Chapter Concepts/Glossary 7-93
 - Before you Begin 7-94
- How Can I Find Out What Happened to An Email Sent to My Blocker? 7-94

Running a Search Query	7-94
Narrowing the Result Set	7-95
Disabling Message Tracking	7-95
How can I View the Activity on My Blocker?	7-95
What Kind of Items Can I Search for in the Email Security Monitor?	7-95
Where Can I View a Snapshot of the Activity on my Blocker?	7-96
System Overview	7-96
Incoming and Outgoing Summary and Graph	7-97
How is Email Categorized?	7-97
Where Can I Find Details about Incoming Mail?	7-98
Notes on Time Ranges in the Mail Trend Graph	7-98
Incoming Mail Details Listing	7-99
Where Can I Find Information about Where My Company Sends Mail?	7-99
Where Can I find Details about the Mail Sent from My Mail Users?	7-99
Where Can I Find Out if My Mail is Being Delivered Correctly?	7-100
Where Can I Drill Down on Information about a Specific Mail User?	7-100
How Can I Find Out Which Content Filters Have Triggered the Most Matches?	7-101
Where Can I Find More Information on the Viruses Detected By My Blocker?	7-101
Where Can I Track Information about My TLS Connections?	7-101
Where Can I Find Out How Well the Blocker Can Handle My Mail Volume?	7-101
Where Can I Find Out More About My System's Health?	7-102
Logging	7-103
Log Rollover and Transfer Schedule	7-103
Log Types	7-103
Retrieving Your Mail Logs	7-105
Troubleshooting Tips	7-106

CHAPTER 8**Advanced Configuration 8-107**

Advanced Configuration: Overview	8-107
Listing of Advanced Topics	8-107

GLOSSARY**APPENDIX A****User License Agreement A-7**

License Agreement	A-7
-------------------	-----



Preface

Revised: October 14, 2009, OL-20859-01

About the Blocker

This chapter describes the organization of the *Cisco Spam & Virus Blocker User Guide* and describes new features available in this release.

You can also view an electronic version of the user guides or request support for your Blocker appliance on the Cisco Support Portal at the following URL:

<http://www.cisco.com/support>

Additional information about the Blocker can be found at:

www.cisco.com/go/blocker

You can visit the site using your support portal account or as a guest user. If you do not already have an account, you can request one.

You might also find it useful to review release notes for earlier releases to see the features and enhancements that were previously added.

What's New in This Release

This section describes the new features and enhancements in the latest Blocker release.

For more information about the current release, see the product release notes, which are available on the Cisco Customer Support Portal.

New features include:

- **Enhanced System Test.** The workflow, functionality, and accessibility are enhanced. See [Enhanced: System Test, page x](#)
- **Configuration Summary Page.** The configuration summary page shows you the settings you have configured for your Blocker. See [New Feature: Configuration Summary Page, page x](#).
- **Categorize Marketing Message as Spam.** You can now categorized marketing messages as spam. See [New feature: Categorize Marketing Messages as Spam, page x](#).
- **Release Notification.** You can now receive notification when new releases are available. See [New feature: Release Notification, page x](#).

New Features and Enhancements

Enhanced: System Test

- **Enhanced testing functionality.** Previously, when you installed your Blocker, a system test was performed to see if the Blocker was prepared to receive mail from external senders, but the system test did not actually send mail to your Blocker. Now, as a part of the system test, Cisco sends a test message from a Cisco server to your Blocker. This ensures that your Blocker effectively handles mail and is ready for use.
- **Enhanced testing workflow.** Previously, if the system test encountered an error, you were required to run the System Setup Wizard again in order to repair any errors. Now, if you encounter an error, the Blocker will attempt to direct you to the part of the configuration that may have errors. When an error occurs, the Blocker redirects you to a part of the configuration file opened for editing. You can change the configuration setting and test the changes immediately.
- **Enhanced accessibility.** Previously, the system test was only accessible via the System Setup Wizard. Now, there is a system test page that you can access at any time. To access this page, go to **System Administration > System Test**. This can simplify your troubleshooting, as the system test is designed to locate common configuration errors.

New Feature: Configuration Summary Page

The configuration summary displays the settings you have configured for your Blocker and allows you to modify these settings from one convenient interface page. It's useful to consult this page if you are having problems sending or receiving mail via your Blocker, as sometimes errors occur in the configuration (for example, a typo exists in the Exchange server name that was entered). It's also a good idea to verify your settings on the Configuration Summary page after you upgrade.

New feature: Categorize Marketing Messages as Spam

Often, we receive marketing messages from companies who have obtained our email addresses through legitimate means (for example, you might get messages from a company where you bought shoes online a few years ago). While this mail is legitimate and not technically, "spam", it may be as much of a nuisance as spam from unsolicited sources. The Blocker can now categorize these marketing messages as spam so you can avoid receiving them.

To enable this feature, ensure that Anti-spam settings are enabled on your default mail policy. Then, go to **Mail Policies > Incoming Mail Policies**. Under the incoming mail policy, you can either click the default mail policy to enable this setting by default, or you can create a new mail policy for a specific group of users. Under the Anti-spam setting, click the section for Anti-spam policies. Scroll down to the **Marketing Email Settings** section, and click **Yes**, to enable marketing email settings.

New feature: Release Notification

As a convenience, Cisco now sends current customers a notification when new releases are ready and prompts customers to upgrade their Blocker appliances. This makes it easy to keep up with the latest releases and bug fixes. To allow Cisco to send you these alerts, go to **System Administration > Alerts**, and click the email address for the user who should receive these alerts. Check the **Release and Support Notifications** checkbox.

How to Use This Guide

Use this guide as a resource to learn about the features and basic configuration of your Blocker appliance. Be aware that *not* all features available on the Blocker appliance are discussed in this guide.

For a list of these features, see the “Advanced Topics” chapter.

The topics in this guide are organized in a logical order, but you may not need to read every chapter in the book. Review the Table of Contents and the section called “How This Book Is Organized” to determine which chapters are relevant to your system.

The guide is distributed electronically as a PDF file. 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.

How This Book is Organized

The following section describes the basic organization and content of the Blocker guide:

- **Installation and Administration of the Blocker.** This chapter includes the steps needed for installation and administration of your Blocker appliance. It includes information about the System Setup Wizard and configuring network settings. It also includes basic administration instructions, such as upgrading, shutting down, and creating users. See [Blocker Installation and Administration, page 1](#).
- **Accepting Mail.** This chapter includes information about configuring your Blocker to accept incoming mail. It describes features you can configure to protect your email infrastructure, such as rate limiting, that prevents your system from being inundated with mail. In addition, this chapter describes the SenderBase reputation service, which uses information about organizations sending mail to you to determine if the mail is spam or not. See [Accepting Mail, page 29](#).
- **Mail Routing for Users and Groups.** This chapter includes information about how to create different mail routing rules for different groups of users. It describes features, such as the Recipient Access Table (RAT) that simplify this process. See [Mail Routing For Users and Groups, page 43](#).
- **Sending Mail.** This chapter includes information about how to configure the Blocker to relay outgoing mail. It describes the process for configuring outbound mail, along with some features that you may want to implement as a part of your company policy. See [Configuring Spam and Virus Protection, page 61](#).
- **Configuring Spam and Virus Protection.** This chapter includes information about configuring Spam and Virus protection. It describes the System Quarantine which quarantines virus-positive emails and emails that are configured for quarantine in the content filters. It also discusses the Spam Quarantine, a special quarantine you can enable to allow your end-users to access their spam messages. Users can also create their own safelists and blocklists to apply to their spam quarantines. See [Configuring Spam and Virus Protection, page 61](#).
- **Creating Custom Rules Using Content Filters.** This chapter describes the content filter tool, which allows you to apply filtering to incoming or outgoing mail based on rules that you created. You can apply these filters for certain groups of users or all users. See [Creating Custom Rules Using Content Filters, page 75](#).
- **Monitoring Your Blocker.** This chapter describes the tools available for monitoring your Blocker and email. Message tracking allows you to track a message as it moves through your email pipeline, while the Email Security Monitor generates different reports related to your incoming and outgoing email. In addition, you can configure log subscriptions to track system health and log system events. See [Monitoring Your Blocker, page 93](#).

- **Advanced Configuration.** This section describes advanced features not discussed in this guide. See [Advanced Configuration, page 107](#).
- **Glossary.** This chapter includes terms you may find useful in administering your Blocker appliance. See [Glossary, page 1](#).
- **User License Agreement.** This chapter includes a detailed user license agreement. See [Advanced Configuration, page 107](#).



CHAPTER 1

Blocker Installation and Administration

Revised: October 8, 2009, OL-20859-01

Blocker Installation and Administration: Overview

This chapter guides you through the process of configuring your Blocker appliance for email delivery using the System Setup Wizard. When you have completed this chapter, the Blocker appliance will be configured to accept mail and protect your network from spam and viruses. This chapter also includes steps for basic administration, such as shutting down your Blocker, adding users, and managing the XML configuration file that stores your configuration settings during upgrades and installations.

Chapter Concepts/Glossary

The following concepts are discussed or introduced in this chapter. Some of these terms are specific to the Blocker appliance, while others are networking terms you will need to know in order to configure your Blocker appliance. It's a good idea to review these terms before you begin.

- **listener.** When you work with the Blocker, you configure what is called a *listener*. A listener's function is to receive messages from either internal mail servers or from external hosts. Once the listener receives mail, it routes it to its eventual destination (either inward or outward); for this reason, listeners are designated as public or private. A *public* listener generally accepts connections from an external host, while a *private* listener generally accepts connections from an internal mail server.
- **SenderBase Reputation Service (SBRS) score.** The SBRS score is a numeric value assigned to an IP address, domain, or organization based on information from the SenderBase Reputation Service. The SenderBase Reputation Service returns a score based on the probability that a message from a given source is spam.
- **Host Access Table (HAT).** The HAT is a table with a list of senders. The HAT applies rules to this list of senders as a way of controlling who can send mail to your network and who can route mail out of your network.
- **Sender Group.** When you create a sender group, you are telling the Blocker to treat a group of users in the same way. A sender group is simply a list of senders gathered together for the purposes of handling email from those senders in the same way. A sender group is a list of senders (identified by IP address, IP range, host/domain, SenderBase Reputation Service classification, SBRS score range, or DNS List query response) separated by commas in a listener's Host Access Table (HAT). For example, you might want to tell the Blocker to accept all mail from senders at *Cisco.com*, but reject all mail from *KnownSpammers.com*.

- **MX record.** An MX record is a record that exists on your DNS server, and it directs the mail flow to your particular domain. Each MX record points to an email server that processes mail for your domain. When you enable the Blocker, you will need to point the MX record so that it directs mail to the Blocker before it can enter your mail server.
- **DNS (Domain Name Service) server.** A domain name system functions as a “phone book” for the internet by translating familiar addresses (such as mycompany.com) into IP addresses (such as 208.77.188.166). A DNS server performs this translation in order to route email to its eventual destination. Your company may have its own DNS server, or it may use an external DNS server. You can perform an MX lookup to see where the DNS server for your domain is.
- **NAT (Network Address Translation).** NAT is the translation of an IP address used within one network to a different IP address used in another network. For example, you might want to route mail to a public IP address while keeping all of your other addresses private. NAT usually occurs on a router or firewall. If your firewall or router is configured for NAT, you’ll need to configure port forwarding to route mail to your Blocker.
- **Remote Host.** A remote host is the sending mail server. For example, if you receive an email from Jenny@Cisco.com, then Cisco.com is the remote host.

Before You Begin

This section describes the steps you should take and the information you should prepare before you start your installation. If you take the time to gather the necessary information and follow these instructions and read the related conceptual information, your installation should go smoothly.

Blocker Pre-Installation Steps

Before you install the Blocker, you need to find and record the following information and keep it handy during the installation.

- **Network settings.** To install the Blocker, you will need to change MX records, firewall, and NAT device settings (if you have NAT devices). Before you perform an installation, find out what your current MX record, firewall, and NAT settings are. If you need more information on MX records, see [DNS Settings, page 1-5](#). For more information about NAT device settings, see [NAT Settings, page 1-6](#). For more information about firewall settings, see [Firewall Settings, page 1-7](#).
- **Settings you will use for the Blocker appliance.** You will need to assign a hostname, IP address, subnet mask, and Gateway IP address to your Blocker appliance. Before you begin installation, be sure you have this information handy.
- **Local DNS Server settings.** If you maintain your own DNS (Domain Name System) server, you will need to obtain the IP address for the primary and secondary DNS servers. Your network administrator should be able to tell you the location of your DNS server if you have one.
- **Mail Configuration settings.** When you set up the Blocker to accept mail, you will need to know the list of domains you will accept email for, and name of the mail server that routes your mail (frequently this is an Exchange server). You should be able to obtain information about mail servers and domains from your mail administrator.
- **Administrator settings.** When you configure the Blocker, you need to designate an administrator for the Blocker appliance. The administrator will receive alerts and mails related to the Blocker appliance and is designated with all Administrator privileges for the Blocker appliance. During the installation, ensure you have the email address for the designated administrator. The initial password for the Blocker is preset and cannot be changed until after the initial installation.

Configuration Options

You can install your Blocker appliance into your existing network infrastructure in several ways. This section addresses some important issues to consider when considering configuration options.

Plan to Place the Blocker Appliance at the Perimeter of Your Network

The Blocker is designed to act as your email “gateway,” receiving and routing the mail that enters your network. Some filtering features, anti-spam and anti-virus features are designed to work with a direct flow of messages from the Internet and from your internal network. The Blocker needs to be accessible via the public Internet as the “first-hop” in your email infrastructure. When the Blocker is the “first hop,” spammers send mail directly to the Blocker appliance, allowing the Blocker to obtain critical connection information about the sender during the SMTP conversation.

The Blocker detects the sender’s IP address and uses this information to determine the reputation score of the sender. The reputation “score” is used to determine if the sender is legitimate or a spammer. In addition, when you place the Blocker at the gateway, the Email Security Monitor feature offers complete visibility into all email traffic for your enterprise from both internal and external senders.

Register the Blocker Appliance in DNS

Malicious email senders actively search public DNS records to hunt for new victims. To fool these senders, you need to register the Blocker in DNS as the primary mail server for your domain. Once the Blocker is registered in DNS, malicious senders will attempt to send mail to the Blocker, allowing the Blocker to gather the connection information it needs to block these senders.

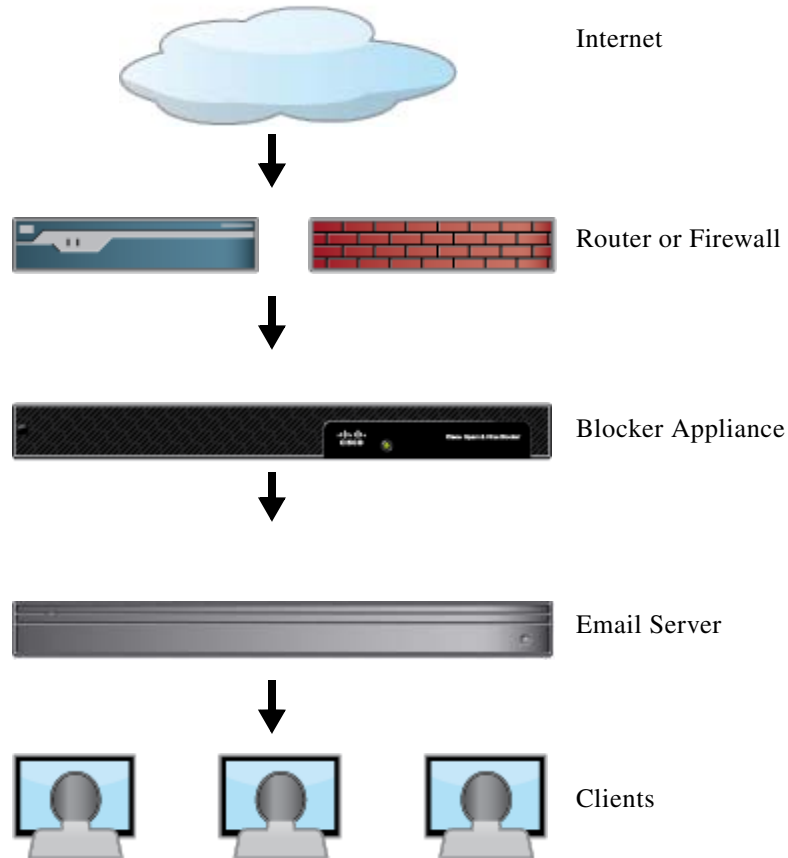
To register the Blocker appliance in DNS, create an A record that maps the Blocker’s hostname to its IP address, and an MX record that maps your public domain to the Blocker’s hostname. You must specify a priority for the MX record to advertise the Blocker appliance as the primary mail server for your domain.

Network Configuration

The following figure shows the general data flow required when routing mail to your Blocker. The important thing to remember is that mail should be routed to the Blocker *first* and then routed to other mail servers.

Figure 1-1 shows the typical placement of the Blocker appliance in a network environment:

Figure 1-1 Network Configuration



Typical Configuration

The typical configuration scenario for the Blocker appliance is as follows:

- **Interfaces** - You configure one or more ethernet interfaces to handle mail. Only one of the two available ethernet interfaces on the Blocker appliance is required for most network environments.
- **Public Listener** (incoming email) - You configure the public listener to accept incoming mail and route it to your incoming mail servers. The public listener receives connections from many external hosts and directs messages to a limited number of internal mail servers.
 - Accepts connections from external mail hosts based on settings in the HAT.
 - 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 mail server, as defined by SMTP Routes.

- **Private Listener** (outgoing email) - You configure a private listener to receive mail from mail servers and send them from your network. The private listener receives connections from a limited number of internal mail servers and directs messages to many external mail hosts.
 - Internal mail servers are configured to route outgoing mail to the Blocker.
 - The Blocker appliance accepts connections from internal mail servers based on settings in the HAT. By default, the HAT is configured to RELAY connections from all internal mail hosts.

DNS Settings

A DNS record is like an entry in an internet “phone book” for your domain. It translates a hostname, such as example.com, into an IP address. Included in the DNS record is an A record that maps the appliance hostname to its IP address and an MX record that directs incoming email to the correct mail server.

If your MX record routes mail to your mail server, you will need to change your MX records to point to your Blocker appliance. If you use a NAT device, you may be able to skip this step (see [NAT Settings, page 1-6](#)).

To change your MX records, locate the MX records on your DNS server. You may have a local DNS server, or your DNS records may be hosted by a DNS provider. The Blocker must be the first hop in your network, so ensure that you configure mail to route through the Blocker before any other mail server.

To change your MX records, consult your DNS administrator or your DNS provider documentation.

In the following example, the MX record pointed to the mail server originally, and is modified to point to the blocker:

Figure 1-2 *Changing Your MX Records***Before**

A Record: exchange.mydomain.com IN A 192.0.2.3

MX Record: mydomain.com in MX exchange.mydomain.com

After

A Record: exchange.mydomain.com IN A 192.0.2.3

A record: mail.mydomain.com IN A 192.0.2.2

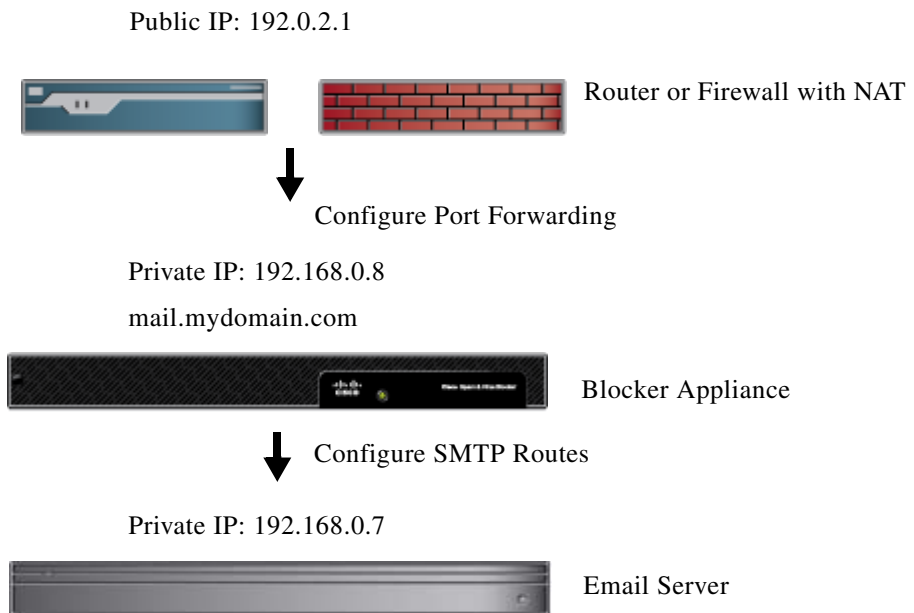
MX record: mydomain.com in MX mail.mydomain.com

NAT Settings

NAT is the translation of an IP address used within one network to a different IP address used in another network. For example, you might want to route mail to a public IP address while keeping all of your other addresses private. If you use Network Address Translation on your router or firewall, you may not need to change your MX records, but you may need to configure port forwarding to ensure mail gets routed to the Blocker.

For instructions on changing your NAT translation tables, consult the documentation for your router or firewall.

In the following example the router/firewall uses NAT to route email from the public IP address of 192.02.1 to the internal IP address of the mail server at 192.168.0.7. The MX records do not need to be modified, but port forwarding must be changed to route port 25 traffic to the Blocker.

Figure 1-3 *Changing Your NAT Settings***Before**

A: mail.mydomain.com IN A 192.0.2.1

MX: mydomain.com IN MX mail.mydomain.com

Port forwarding: Port 25 traffic to 192.168.0.7

After

A: mail.mydomain.com IN A 192.0.2.1

MX: mydomain.com IN MX mail.mydomain.com

Port forwarding: Port 25 traffic to 192.168.0.8

SMTP route between Blocker and mail server

Firewall Settings

The following table lists the possible ports that may need to be opened for proper operation of the Blocker appliance (these are the default values).

Port	Protocol	In/Out	Hostname	Description
20/21	TCP	In or Out	Blocker IPs, FTP Server	FTP for aggregation of log files.
22	TCP	In	Blocker IPs	SSH access, 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	Blocker IPs	Telnet access, 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	Blocker IPs	SMTP to receive bounced email or if injecting email from outside firewall.
80	HTTP	In	Blocker IPs	HTTP access to the GUI for system monitoring.
80	HTTP	Out	downloads.ironport.com	Service updates
80	HTTP	Out	updates.ironport.com	Blocker upgrades
82	HTTP	In	Blocker IPs	Used for viewing the Blocker Spam quarantine.
83	HTTPS	In	Blocker IPs	Used for viewing the Blocker Spam quarantine.
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 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 Spam Quarantine
161	UDP	In	Blocker 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 Spam Quarantine
636 3269	LDAPS	Out	LDAPS	LDAPS — ActiveDirectory's Global Catalog Server
443	TCP	In	Blocker IPs	Secure HTTP (https) access to the GUI for system monitoring.
443	TCP	Out	update-static.ironport.com	Verify the latest files for the update server.
514	UDP/TCP	Out	Syslog Server	Syslog logging
628	TCP	In	Blocker IPs	QMQP if injecting email from outside firewall.
6025	TCP	Out	Blocker IPs	Spam Quarantine

Running the System Setup Wizard

The Blocker provides a browser-based System Setup Wizard to guide you through the process of system configuration. If you have gathered the information required in [Blocker Pre-Installation Steps, page 1-2](#) the configuration process will take less time to complete.

The Blocker appliance ships with a default IP address of 192.168.42.42 on the Data 1 port of appliance. Before connecting the Blocker appliance to your network, ensure that no other device's IP address conflicts with this factory default setting.



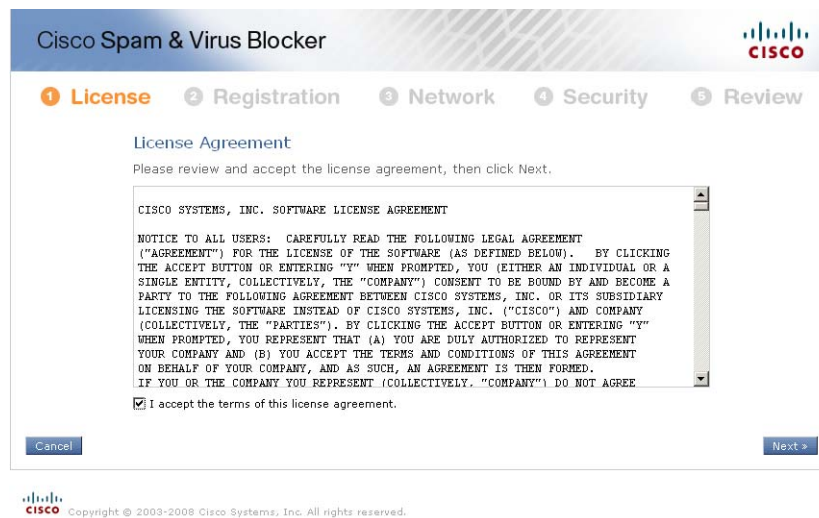
Warning

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.

To run the System Setup Wizard, access the System Setup Wizard from **System Administration > System Setup Wizard**.

Step 1: Accept license agreement.

Figure 1-4 The System Setup Wizard - Accept License Page



Review the license agreement, select **accept the terms of this license agreement**, and click **Next**.

Step 2: Enter registration information.**Figure 1-5 The System Setup Wizard - Registration Page**

Enter your company and contact information, and click **Next**.

Step 3: Enter network information.**Figure 1-6 The System Setup Wizard - Network Page**

Section	Field	What You Need to Enter
Network Settings	Blocker Hostname	Provide a fully qualified hostname for Blocker so you can access it later. Example: <i>blocker.mycompany.com</i>
	Blocker IP Address	Provide the IP address you'd like to assign to the Blocker. The Blocker port labeled "Data 2" will be assigned this IP address. Example: <i>10.251.21.11</i>
	Subnet Mask	Provide a subnet mask to define the segment of the network the Blocker belongs to. Example: <i>255.255.255.0</i>
	Gateway IP Address	Enter the gateway address of the router which the Blocker communicates with to send and receive data from the internet. Example: <i>10.251.21.1</i>
	Time Zone	Your time zone selection automatically updates for Daylight Savings.
Mail Configuration	DNS	Blocker can use the Internet root DNS servers or your own DNS servers that you specify. Custom DNS example: <i>10.251.21.15</i>
	Help Fight Spam	Share information about your mail flow to help Cisco fight spam. The information shared is not personally identifiable and is used only to help spam fighting engines work more effectively.
	Accept Mail for these Domains	Enter domains(s) you would like Blocker to accept mail for. Separate multiple domains with commas. Example: <i>mycompany.com</i>
Administrator Settings	Exchange/Mail Server	Enter a hostname or IP address for your Exchange/Mail server so that the Blocker knows where to send mail. Example: <i>exchange.mycompany.com, 10.251.21.15</i>
	Administrator Email	The administrator email will be used to send scheduled reports and system alerts and feature key renewal notices. Example: <i>admin@mycompany.com</i>
	Your new Administrator Password	Your password must be 6 or more characters. Please make a note of your new password. It will take effect after the System Setup Wizard is complete. The default administrator username is "admin" and is not editable.

Step 4: Set anti-spam and anti-virus policy and install the configuration.**Figure 1-7 The System Setup Wizard - Security Page**

Cisco Spam & Virus Blocker

1 License 2 Registration 3 Network 4 Security 5 Review

Anti-Spam Policy

How would you like to handle spam?

- Block and Quarantine (recommended)
 - and notify users with a daily digest.
- Quarantine All
 - and notify users with a daily digest.
- Tag and Deliver

Anti-Virus Policy:

- Block Viruses

More Information

Mouse over fields to the left for additional notes.

Anti-Virus technology protects against multiple threats like worms, Trojans, viruses.

< Previous Cancel Next >

Configure anti-spam and anti-virus settings.

Step 5: Review and install the configuration.**Figure 1-8 The System Setup Wizard - Review Page**

Cisco Spam & Virus Blocker

1 License 2 Registration 3 Network 4 Security 5 Review

Review Your Configuration

Please review your configuration. If you need to make changes, click the edit link to return to the page you'd like to edit.

Registration [Edit]

Company Information: MyCompany
123 Company Lane
Company City, CA 94117
United States

Primary Contact: Jane Doe
jdoe@example.com
Business: 555 555 5555
Mobile: 555 555 5555

Secondary Contact: N/A

Network [Edit]

Blocker Hostname: blocker.example.com

Blocker IP Address (Data 2): 192.168.42.43
Subnet Mask: 255.255.255.0

Gateway IP Address: 192.168.42.42
Time Zone: Europe/Dublin

DNS: Use the Internet's Root DNS servers

Help Fight Spam: Disabled

Accept Mail for these domains: mycompany.com

Exchange/Mail Server: exchange.mycompany.com

Administrator Email: jdoe@example.com

Administrator Password: (hidden)

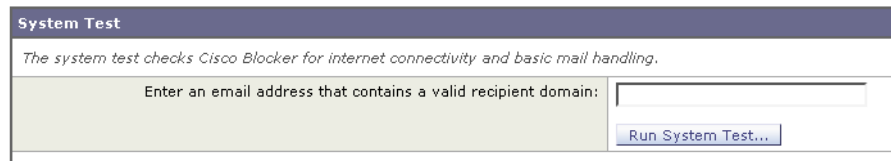
Review your configuration and accept it when you are satisfied that you filled out all the fields correctly.

Step 6: Run the System Test

Log back into the appliance with the password you set in the System Setup Wizard. Once you log back into the appliance, the GUI opens to the “Next Steps” page which directs you to run a system test.

Figure 1-9 The System Test Page - Enter Email Address

System Test



The screenshot shows a web interface for the 'System Test' page. At the top, there is a header 'System Test' and a subtitle 'The system test checks Cisco Blocker for internet connectivity and basic mail handling.' Below the subtitle is a text input field with the placeholder text 'Enter an email address that contains a valid recipient domain:' and a 'Run System Test...' button.

To run the system test, enter a valid email address that is in one of the domains you configured your Blocker to accept mail for. The system test will attempt to send an email to that address. If problems occur during the system test, the System Setup wizard will redirect you to the Configuration Summary page and allow you to change the settings that are likely to be causing your problems.

The system test performs four basic steps:

- **Verifies internet connection.** First, the system tests the connection between the Blocker and the internet.
- **Connects to your mail server.** Next, the Blocker tests the connection between the mail server you configured and the Blocker.
- **Verifies that your Blocker can receive mail.** After connecting to the mail server, the Blocker verifies that it can receive a message.
- **Sends a test message.** Once all these systems are tested, a server at Cisco sends a test message to your Blocker appliance. You should be able to retrieve the test message at the email address you entered in the “System Test” page.



Note

To ensure you receive your test mail, you may need to configure your Outlook filters to skip automatic filtering. To do this, go to **Actions > Junk Email > Junk Email Options > Options** page. Set the level of junk email protection to "No automatic filtering."

Figure 1-10 The System Test Page - Test Progressing

System Test

The system test checks Cisco Blocker for internet connectivity and basic mail handling.

Enter an email address that contains a valid recipient domain:

✔ **Verifying internet connection...**

✘ **Connecting to your Exchange/Mail Server...**

Blocker was unable to establish a connection with your Exchange/Mail Server mail.qa

There was a problem connecting to your Exchange/Mail server with Blocker's current setting of port 25. This is usually caused by one of the following problems:

- Your Exchange/Mail service is not up.
- Your Exchange/Mail Server is not configured to accept mail from this address.
- Blocker is misconfigured with an incorrect hostname for the Exchange/Mail Server.
 - > [Check your Exchange/Mail Server hostname settings.](#)

Verifying that Blocker can receive mail...

Sending a test message...

System test complete.

As the test progresses, a green check indicates that the section of the test passed. If a red “X” appears, you need to troubleshoot a problem in your configuration and run the test again. When possible, the Blocker will identify problems and link you to the place where you can fix the issue. Clicking the link will take you to the section of the Configuration Summary page which the error indicates needs to be modified.

For example, in the previous screen, the test shows a problem connecting to the Exchange server. If you click on the link, the Configuration Summary page opens, allowing you to edit the setting.

Figure 1-11 The System Test Page - Changing Settings

Troubleshoot — Blocker is unable to contact your Exchange/Mail Server using the current hostname. Review the settings below to determine if there is a configuration problem and make any needed changes.

Configuration Summary					
Network					
Incoming Listener:	IncomingMail				
Blocker Hostname:	blocker03.qa				
Blocker IP Address (Incoming Mail):	1.1.1.1 (Data 2)				
Subnet Mask:	255.255.255.0				
Default Gateway IP Address:	10.92.145.1				
Time Zone:	(GMT) Greenwich Mean Time (Dublin, Edinburgh, Lisbon, London)				
DNS Servers:	10.92.144.4				
Accept Mail for Domain:	qa47.qa				
Exchange/Mail Server:	<table border="1"> <thead> <tr> <th>Server</th> <th>Domains</th> </tr> </thead> <tbody> <tr> <td>example.qa</td> <td>example1.qa</td> </tr> </tbody> </table> <p><i>Is this the correct information for your Exchange or mail server? Make any needed corrections and save your changes.</i></p> <p><i>Enter a single server hostname or IP address, and one or more domains, for each Exchange server. Use commas to separate entries if you enter multiple domains per exchange server. Examples: exchange.mycompany.com; .myco, mycompany.com, internal-mycompany.com.</i></p>	Server	Domains	example.qa	example1.qa
Server	Domains				
example.qa	example1.qa				
Alert Recipients:	test@mail.qa				
Security Services					
Default Incoming Anti-Spam Policy:	Positive: Block Suspect: Quarantine				
Default Incoming Anti-Virus Policy:	Enabled				
Help Fight Spam with Senderbase:	Enabled				

Cancel Save/Test

**Note**

When you open the Configuration Summary page from the link in the System test, the page opens in “Troubleshooting mode” and you can only edit the sections that testing errors indicate are a problem. If you want to edit the entire configuration summary, go to **System Administration > Configuration Summary**, and click **Edit** to open the entire page for editing.

Once you edit the settings, click the **Save/Test** button. On most Blocker pages, you have to commit your changes to save them, but in this case, the changes are saved in a single click.

There are a few problems that commonly occur when configuring the Blocker. You may need to change your MX records to point to the Blocker appliance. It can take some time for the MX record changes to take place, so verify that you have given enough time for the MX records to change over. Other common connectivity issues include problems with the firewall (you need to open ports 25, 80, and 443). In addition, port 25 traffic must be directed to the Blocker. You may also need to configure port forwarding if you use a NAT device.

**Note**

You can run the system test at any time to verify your configuration. The system test is located at **System Administration > System test**.

As a final check, you can go to **Monitor > Incoming Mail** to view your incoming mail reports. Your mail delivery should be displayed.

Configuring Your System to Use LDAP Authentication

If you use LDAP authentication in your corporate infrastructure, you can also leverage your LDAP directories to work with your Blocker appliance. This means you don’t have to recreate your mail user list on your Blocker, and you can go to a single place to authentic your users.

If you are running an Active Directory server on your network, you can use the Active Directory Wizard to configure LDAP settings. Because LDAP syntax varies widely, it's difficult to configure these settings manually. Using the Active Directory wizard automates this process and prevents potential errors in creating the most commonly used LDAP queries: ACCEPT and GROUP queries. Other LDAP directories are supported on your Blocker, but you cannot use the Active Directory Wizard to set up authentication. For instruction on setting up other types of LDAP Authentication, see the *AsyncOS for Email Advanced User Guide* (PDF available on your Documentation CD).

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.

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.

Running the Active Directory Wizard:

1. On the Active Directory Wizard page, click **Run Active Directory Wizard**.
2. Enter the host name for the Active Directory server.
3. Enter a username and password for the authentication request.
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.

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.
6. Click **Done**.

The Blocker Interface

Logging In

The first time you log into the Blocker appliance, you will log onto a factory-configured address. During the installation, you can change this address.

To access the GUI, open your web browser and point it to **http://192.168.42.42**.

The login screen is displayed:

Log into the appliance by entering the factory default username and password below.

- Username: **admin**
- Password: **cisco**

Your session will time out if it is idle for over 30 minutes or if you close your browser without logging out. If this happens, 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.

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 Blocker appliance:

- **Monitor.** This tab includes features that allow you to monitor your email traffic, including reports, quarantines, and message tracking features.
- **Mail Policies.** These “policies” allow you to have granular control over the mail entering and exiting your network. You can create a “policy” to control different aspects of mail delivery.
- **Security Services.** This tab includes security features, such as anti-spam and anti-virus.
- **Network.** This tab includes features to configure different network configuration settings, such as DNS, listeners, and IP interfaces.
- **System Administration.** This tab includes features having to do with basic system administration such as shutting down, rebooting, setting alerts, adding users, testing your configuration, and viewing your system configuration settings, etc.

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 hovering over the menu headings for a main section (Monitor, Mail Policies, Security Services, Network, and System Administration) and navigating to the appropriate subsection.

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**.

Saving Your Changes

As you make configuration changes in the GUI, you *must* explicitly commit those changes by clicking the **Commit Changes** button. If you do not commit changes, your work will not be saved.

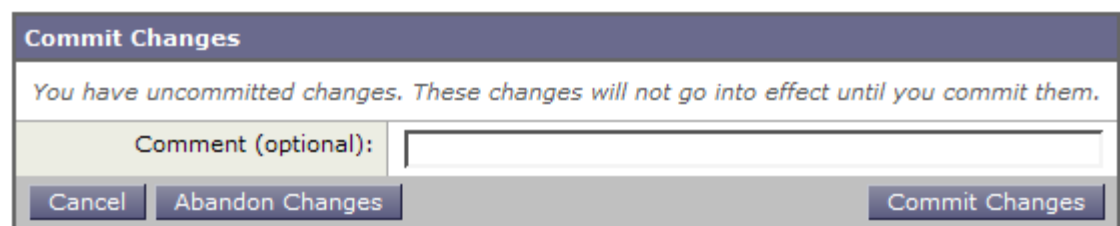
Figure 1-12 Commit Changes Button



Clicking the **Commit Changes** button displays a page where you can add a comment and commit the changes, or abandon all changes made since the most recent commit.

Figure 1-13 Commit Changes Page

Uncommitted Changes



Throughout this guide, you will see instructions that tell you to “submit and commit your changes.” When you see this instruction, you’ll need to click the **Commit Changes** button as shown above.

The Configuration Summary Page

The configuration summary page allows you to review and edit the configured settings for your Blocker at a glance. You can edit the configured settings by clicking the **Edit** button at the bottom of the page. This is particularly useful when you are troubleshooting an installation, about to perform an upgrade, or performing post-upgrade troubleshooting.

Figure 1-14 The Configuration Summary Page(Viewing Mode)
Configuration Summary

Configuration Summary	
Network	
Incoming Listener:	IncomingMail
Blocker Hostname:	blocker03.qa
Blocker IP Address (Incoming Mail):	1.1.1.1 (Data 2)
Subnet Mask:	255.255.255.0
Default Gateway IP Address:	10.92.145.1
Time Zone:	(GMT) Greenwich Mean Time (Dublin, Edinburgh, Lisbon, London)
DNS Servers:	10.92.144.4
Accept Mail for Domain:	qa47.qa
Exchange/Mail Server:	mail.qa
Alert Recipients:	admin@blocker03.qa
Security Services	
Default Incoming Anti-Spam Policy:	Positive: Block Suspect: Quarantine
Default Incoming Anti-Virus Policy:	Enabled
Help Fight Spam with Senderbase:	Enabled
Edit Settings...	

When you edit the configuration summary file, you can *change* settings, but you cannot *add* settings. For example, if you want to add a second listener, you must go to the **Network > Listeners** page to do so.

When you click **Edit**, the page opens to allow you to edit your current settings:

Figure 1-15 *Editing the Configuration Summary Page 1 (Edit Mode)*
Configuration Summary

Configuration Summary	
Network	
Incoming Listener:	IncomingMail
Blocker Hostname:	blocker03.qa
Blocker IP Address (Incoming Mail):	1.1.1.1 (Data 2) <small>Enter an IP address for the incoming mail data port. Example: 10.251.21.15.</small>
Subnet Mask:	255.255.255.0 <small>Enter a value for the subnet mask for the incoming mail data port. Example: 255.255.255.0.</small>
Default Gateway IP Address:	10.92.145.1 <small>Enter a value for the default gateway IP address</small>
Time Zone:	(GMT) Greenwich Mean Time (Dublin, Edinburgh, Lisbon, London) ▼
DNS Servers:	<input type="radio"/> Use the Internet's Root DNS Servers <input checked="" type="radio"/> Use the specified DNS Servers: <input type="text" value="10.92.144.4"/> <small>Enter an IP address for each DNS server. Enter only one IP address per field. Example: 10.251.21.5.</small>
Accept Mail for Domain:	qa47.qa <small>Enter one hostname or IP address for each domain for which Blocker will accept mail. Enter only one item per field. Examples: mycompany.com, 10.251.21.15</small>
Exchange/Mail Server:	Server: <input type="text" value="mail.qa"/> Domains: <input type="text" value="qa47.qa"/> <small>Enter a single server hostname or IP address, and one or more domains, for each Exchange server. Use commas to separate entries if you enter multiple domains per exchange server. Examples: exchange.mycompany.com; .myco, mycompany.com, internal-mycompany.com.</small>

Figure 1-16 *Editing the Configuration Summary Page 2 (Edit Mode)*

Alert Recipients:	<input type="text" value="test@mail.qa"/> <small>Enter a valid email address for each intended alert recipient. Enter only one email address per field. Example: admin@mycompany.com</small>
Security Services	
Default Incoming Anti-Spam Policy:	Positively Identified Incoming Spam <input checked="" type="radio"/> Block <input type="radio"/> Quarantine <input checked="" type="checkbox"/> and notify users with a daily digest. <input type="radio"/> Tag and Deliver Suspected Incoming Spam <input type="radio"/> Block <input checked="" type="radio"/> Quarantine <input checked="" type="checkbox"/> and notify users with a daily digest. <input type="radio"/> Tag and Deliver
Default Incoming Anti-Virus Policy:	<input checked="" type="checkbox"/> Enable
Help Fight Spam with SenderBase:	<input checked="" type="checkbox"/> Enable sharing limited data with the SenderBase Information Service (Recommended)
<input type="button" value="Cancel"/> <input type="button" value="Save/Test"/>	

After you make changes, click the **Save/Test** button.

The Configuration Summary Page:

Field	Description
Incoming Listener	Name of the configured public listener. Example: <i>PublicListener1</i>
Blocker Hostname	The fully qualified hostname you assigned to the Blocker. Example: <i>blocker.mycompany.com</i>
Blocker IP Address (Incoming Mail)	The IP address assigned to the Blocker. The Blocker port labeled "Data 2" is assigned to this IP address by default. Example: <i>10.251.21.11</i>
Subnet Mask	The subnet mask that defines the segment of the network the Blocker belongs to. Example: <i>255.255.255.0</i>
Default Gateway IP	The gateway address of the router which the Blocker communicates with to send and receive data from the internet. Example: <i>10.251.21.1</i>
DNS Server	The local DNS server you specified during your Blocker installation (if applicable). Custom DNS example: <i>10.251.21.15</i>
Accept Mail for Domain	Domain or domains you configured the Blocker to accept mail for. Example: <i>mycompany.com</i>
Exchange/Mail Server	The hostname or IP address for your Exchange/Mail server associated with your Blocker. Example: <i>exchange.mycompany.com, 10.251.21.15</i>
Alert Recipients	The email address configured for receiving email alerts.
Default Incoming Anti-Spam Policy	The basic configuration of the incoming spam policy. Examples: <ul style="list-style-type: none"> • <i>Block & Quarantine</i> • <i>Quarantine All</i> • <i>Tag & Deliver</i>
Default Incoming Anti-Virus Policy	The basic configuration of the incoming virus policy Example: <i>Enabled</i>
Help Fight Spam with Senderbase	Whether sharing information with Senderbase is enabled or disabled.

System Administration

You can perform most administrative tasks from the System Administration menu in the Graphical User Interface (GUI).

Shutting Down the Blocker

To safely power down your Blocker appliance, go to **System Administration > Shutdown/Reboot** and choose shutdown as the operation to perform. Select the number of seconds before the action will take place, and click **Commit**.

**Note**

It is important that you shut down your Blocker using the shutdown page rather than performing a hard reboot. This protects your files and prevents corruption of your queue.

You can restart the appliance later without losing any messages queued up for delivery. The default delay is thirty (30) seconds. The Blocker allows open connections to complete during the delay, after which it forcefully closes open connections.

Rebooting the Blocker

To safely reboot your Blocker, go to **System Administration > Shutdown/Reboot** and choose reboot as the operation to perform. Select the number of seconds before the action will take place, and click **Commit**.

The Blocker reboots itself safely without losing any messages queued up for delivery. The default delay is thirty (30) seconds. The Blocker allows open connections to complete during the delay, after which it forcefully closes open connections.

Open a Support Case

To open a technical support case, go to **Help and Support > Open a Support Case**.

The Support case page opens:

Figure 1-17 Open a Support Case

Open a Technical Support Case

Technical Support Case

Send Request to:
Separate multiple email addresses with commas.

Contact Information:

Name:

Email:

Other Contact Information (optional)

Phone1:

Phone2:
(Mobile, Pager, etc.)

Other:

Issue Description:

Issue Subject:

Issue Description:

Send

Enter information in the fields and click **Send**.

Upgrading the Blocker

Before you upgrade your Blocker, you need to configure your update settings to tell your Blocker where and how to pick up the upgrade files. You can configure your Blocker to get upgrades from a local server or from a Cisco server. Select the interface to use for the upgrade, and configure proxy settings if you wish. For details, see [Editing Update Settings, page 1-23](#).

Use the following guidelines to upgrade your Blocker:

1. Export the XML configuration file off the Blocker.
2. If you are using the Safelist/Blocklist feature, export the list off the Blocker.
3. Go to the **System Administration > System Upgrade** page.
4. Click **Available Upgrades**.
5. Select an upgrade from the list of available upgrades.
6. If you haven't already exported your XML configuration file, select the option to save the XML configuration file to the Configuration directory and email it off-box.
7. 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.
8. After the upgrade is finished, you are asked to reboot the appliance.
9. Click **Reboot Now**.

Working with Feature Keys

Occasionally, your support team may provide a key to enable specific functionality on your system. Use the **System Administration > Feature Keys** page in the GUI to enter the key and enable the associated functionality.

Keys are specific to the serial number of your appliance and specific to the feature being enabled (you cannot re-use a key from one system on another system). If you incorrectly enter a key, an error message is generated.

Feature keys functionality is split into two pages: Feature Keys and Feature Key Settings.

The Feature Keys Page

Log in to the GUI and click **System Administration > Feature Keys**.

The Feature Keys page:

- Lists all active feature keys for the appliance
- Shows any feature keys that are pending activation
- Looks for new keys that have been issued (optional, and also can install keys)

A list of the currently enabled features is displayed. The Pending Activation section is a list of feature keys that have been issued for the appliance but have not yet been activated. Your appliance may check periodically for new keys depending on your configuration. You can click the **Check for New Keys** button to refresh the list of pending keys.

Feature Key Settings

The Feature Key Settings page is used to control whether your appliance checks for and downloads new feature keys, and whether or not those keys are automatically activated.

To add a new feature key manually, paste or type the key into the Feature Key field and click **Submit Key**. An error message is displayed if the feature is not added (if the key is incorrect, etc.), otherwise the feature key is added to the display.

To activate a new feature key from the Pending Activation list, select the key (mark the “Select” checkbox) and click **Activate Selected Keys**.

You can configure your Blocker to automatically download and install new keys as they are issued. In this case, the Pending Activation list will always be empty. You can enable the Blocker to check for new keys at any time by clicking the **Check for New Keys** button, even if you have disabled the automatic checking via the Feature Key Settings page.

Expired Feature Keys

If the feature key for the feature you are trying to access has expired, please contact Cisco Customer Support.

Editing Update Settings

Many of the settings used to configure how the Blocker updates various services (such as the anti-spam and anti-virus services) are accessible via the Service Updates page from the Security Services menu.

The Service Updates Page

The Service Updates page displays the current settings for updating various services for your Blocker. To access the Service Updates page, go to **Security Services > Service Updates**.

The update servers use dynamic IP addresses. If you have strict firewall policies, you may need to configure a static location for updates. If you determine that your firewall settings require a static IP address for updates, follow instructions below for editing the update settings and contact Cisco Customer support to obtain the required URL addresses.

Editing Update Settings

To edit the update settings for your Blocker, click the **Edit Update Settings** button. The Edit Update Settings page is displayed. On the Edit Update Settings page, you can edit the following options:

- Automatic updates. Enable automatic updates and the polling interval (how often the appliance will check for update) for Sophos Anti-Virus definitions and IronPort Anti-Spam rules.
- **Update Intervals.** Determine the intervals for automatic updates. Use a trailing 'm' for minutes, 'h' for hours or 'd' for days. The minimum valid update time is 5m or enter '0' to disable automatic updates (manual updates will still be available for individual services).
- **Interface** (Applies only to Blocker upgrades). Select a network interface to accept Blocker upgrades. By default, the appliance selects an interface to use.
- **HTTP Proxy Server.** An optional proxy server used for all of the following services:
 - Sophos Anti-Virus definitions
 - IronPort Anti-Spam rules
 - Feature Key updates
 - Blocker upgrades



Note Note that if you specify a proxy server, it will be used for all of the services listed above.

- **HTTPS Proxy Server.** An optional proxy server using HTTPS. If you define the HTTPS proxy server, it will be used to update the following services:
 - Symantec Brightmail Anti-Spam rules
 - Blocker upgrades
 - SenderBase Network Participation sharing

Adding Users

You can add users in two ways:

- create a user account on the Blocker appliance
- enable the Blocker to use your company's user authentication directory (such as an LDAP directory). This allows you to access your company's email addresses without having to manually add each user.

You can manage users and connections to external authentication sources on the **System Administration > Users** page in the GUI.

For each new user account you create on the Blocker appliance, you specify a username and a full name, and then assign the user to one of the following user roles:

- **Administrator.** User accounts with the Administrator role have full access to all configuration settings of the system. The admin user account cannot be edited or deleted, aside from changing the password. To change the password for the default admin user account, use the Edit User page in the GUI. If you forget the password for the admin user account, contact your customer support provider to reset the password.
- **Operator.** User accounts with the Operator role are restricted from: Creating or editing user accounts, performing some quarantine functions (including creating and deleting quarantines), running the System Setup wizard, and resetting the configuration.
- **Guest.** Users accounts with the Guest role can only view status information. Users with the Guest role can also manage messages in the IronPort Spam Quarantine and system quarantines, if access is enabled. Users with the Guest role cannot access Message Tracking.
- **Read-Only Operator.** User accounts with the Read-Only Operator role have access to view configuration information. Users with the Read-Only Operator role can make and submit changes to see how to configure a feature, but they cannot commit them. Users with this role can manage messages in the IronPort Spam Quarantine and system quarantines, if access is enabled. Users with this role cannot access the file system, FTP, or SCP.
- **Help Desk User.** User accounts with the Help Desk User role are restricted to: Message Tracking and managing the IronPort Spam Quarantine and system quarantines.

Each role contains differing levels of permissions within the system. After you have assigned a role, you specify a password for the user.

The default user account for the system, admin, has all administrative privileges.

Changing Your Password

Users can change their own passwords via the **Options > Change Password** link at the top of the GUI.

Enter the old password then enter the new password and retype it for confirmation. Click **Submit**. You are logged out and taken to the log in screen.

Configuring the Return Address for Various Generated Messages

You can configure the envelope sender for mail generated by Blocker 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.

Alerts

Alerts are email notifications containing information about events occurring on the Blocker 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 Blocker 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. To edit alert settings, click **Edit Settings** on the Alerts page.

You also use alerts to enable Cisco to send you notification of new releases and upgrades. To allow Cisco to send you these alerts, go to **System Administration > Alerts**, and click the email address for the user who should receive these alerts. Check the **Release and Support Notifications** checkbox.

Configuring Domain Name System (DNS) Settings

You can configure the DNS settings for your Blocker appliance through the DNS page on the Network menu of the GUI.

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

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.

System Time

To set the System Time on your Blocker 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

Managing the Configuration File

All configuration settings within the Blocker can be managed via a single configuration file. The file is maintained in XML (Extensible Markup Language) format.

You can use this file in several ways:

- You can save the configuration file to a different system to back up and preserve crucial configuration data. If you make a mistake while configuring your appliance, you can “roll back” to the most recently saved configuration file.
- You can download the existing configuration file to view the entire configuration for an appliance quickly. (Many newer browsers include the ability to render XML files directly.) This may help you troubleshoot minor errors (like typographic errors) that may exist in the current configuration.

- You can download an existing configuration file, make changes to it, and upload it to the same appliance. This, in effect, “bypasses” the GUI for making configuration changes.
- You can upload entire configuration file via FTP access.
- Because the file is in XML format, an associated DTD (document type definition) that describes all of the XML entities in the configuration file is also provided. You can download the DTD to validate an XML configuration file before uploading it. (XML Validation tools are readily available on the Internet.)

Troubleshooting Tips

This section addresses common problems that occur when installing the Blocker:

- **Problem:** Why don't I see messages in the mailbox for a newly added domain, even though the Incoming Mail report shows that hundreds of messages have been processed?
- **Solution:** You can troubleshoot delivery status by reviewing delivery status details from the **Monitor > Delivery Status** page. A common cause of this problem is that an SMTP route is missing for your mail server. When you add a new domain to the Recipient Access Table (RAT), you also need to add a corresponding SMTP route on the **Network > SMTP Route** page. The Blocker cannot deliver mail to users until an SMTP route exists for the mail server.
- **Problem:** I ran the System Setup Wizard without network connectivity. How do I know if the Blocker can successfully deliver mail?
- **Solution:** After connecting the Blocker to the network, review your system configuration settings and verify that you have configured MX records, NAT settings, and any necessary firewall settings. Then, go to **System Administration > System test** and run the system test to ensure you have configured Blocker settings correctly.



CHAPTER 2

Accepting Mail

Revised: October 8, 2009, OL-20859-01

Accepting Mail: Overview

After you have configured the basic settings in your Blocker via the System Setup Wizard, you are now ready to begin tailoring the way your Blocker will handle incoming mail. All mail routed to the Blocker is routed through a **listener**. A listener is an email processing service that receives and delivers email on a particular IP interface.

You create a set of rules for incoming mail that you can then apply to a specific listener--this allows you the flexibility to create different rules for each listener.

One of the most powerful tools for controlling and filtering incoming mail is the *Host Access Table*, or HAT. The HAT is a table you create that contains a list of senders. You then apply rules to this list of senders as a way of controlling who can send mail to your network and who can send mail out of your network.

You also configure SBRS (Senderbase Reputation Score) settings. An SBRS score 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. When you configure this setting, you can control the amount of suspect spam that enters your network.

Chapter Concepts/Glossary

The following concepts are discussed or introduced in this chapter. Some of these terms are specific to the Blocker appliance, while others are networking terms you will need to know in order to configure your Blocker appliance. It's a good idea to review these terms before you begin.

- **listener**. When you work with the Blocker, you configure what is called a **listener**. A listener's function is to receive messages from either internal mail servers or from external hosts. Once the listener receives mail, it is routed to its eventual destination (either to a user's inbox or out towards a destination mailbox); for this reason, listeners are designated as public or private. A *public* listener generally accepts connections from a sending mail server (sometimes called a *remote host*), while a *private* listener generally accepts connections from an internal mail server.
- **Host Access Table (HAT)**. The HAT is a table with a list of senders. The HAT applies rules to this list of senders as is a way of controlling who can send mail to your network and who can route mail out of your network.

- **Remote host.** The remote host is usually a mail server trying to connect to your network in order to deliver mail. You will define some rules for the sender or remote host to determine how to handle the mail that is received from this host.
- **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. A sender group is a list of senders (identified by IP address, IP range, host/domain, SenderBase Reputation Service classification, SBRS score range, or DNS List query response) separated by commas in a listener's Host Access Table (HAT).
- **Mail Flow Policy.** A mail flow policy is a way of grouping parameters (an access rule, followed by rate limiting parameters and custom SMTP codes and responses) in the Host Access Table (HAT) for a listener. Together, sender groups and mail flow policies are defined in a listener's HAT. The Blocker ships with some predefined mail flow policies and sender groups for listeners.
- **SenderBase Reputation Service (SBRS) score.** An SBRS score 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

Before You Begin

Before you begin configuring settings for mail delivery, you should have run the system setup wizard, which configures a few basic settings for the Blocker:

- Management interface on the Data 1 port
- Incoming mail configured on your Data 2 port
- Opened appropriate ports on your firewall
- Modified your MX records to route mail to the Blocker
- Configured NAT settings for port forwarding (if you have NAT configured).

If you have configured all these settings, you are ready to configure other mail settings to control how your Blocker accepts mail.

Understanding How the Blocker Accepts Mail (Public Listeners)

When you ran the System Setup Wizard, you configured the Blocker to accept mail on a public listener. This means that you configured the Blocker to *receive* mail.

When you configured your Blocker for mail delivery, you enabled mail acceptance on a *listener*. A listener describes an email processing service that is configured on a particular IP interface. Listeners process mail entering the Blocker — either from the internal systems within your network or externally from the Internet. You can think of a listener as an “email injector” running on a specific port for each IP address you specify.

To configure listeners manually, go to Listeners page (**Network > Listeners**). When you create a new listener via the System Setup Wizard, the listener is configured with default SBRS (SenderBase Reputation Service) values. However, when you create a listener manually, it is not configured with these default values.

**Note**

You cannot configure mail to be delivered on multiple ports on a single IP address (for example, port 25 for normal delivery and port 6025 for the spam quarantine). Cisco recommends running each delivery option on a separate IP address or host. Also, you cannot use the same hostname for regular email delivery and quarantine delivery.

Sending and Receiving Mail on a “Listener”

The Blocker differentiates between *public* listeners — which are intended for receiving email from the Internet — and *private* listeners that are intended to accept email only from internal mail servers. By creating distinct public and private listeners for different public and private networks, you can configure different filtering options for inbound and outbound mail. 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.

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 Spam Quarantine via an IP interface.

What if My Blocker Isn't at the Edge of My Network?

In some rare cases, you may need to maintain a mail server at the edge of your network (rather than placing it after the Blocker in the network). This can pose a problem for the Blocker because the Blocker relies on information received from the remote host (sending mail server) to secure your network. IronPort Anti-Spam depends on accurate IP addresses for external senders so it is vital for the Blocker to have this information. To ensure that the Blocker can get accurate information in this configuration, you need to set up an incoming relay. The Incoming Relays feature helps your Blocker obtain the IP address of an external machine that is sending mail to the Blocker via one or more mail exchange/transfer agents (MX or MTA) at the edge of the network.

When configuring an incoming relay, you specify the names and IP addresses of all of the internal mail servers connecting to the Blocker appliance, as well as the header used to store the originating IP address.

Enabling the Incoming Relays Feature

Once enabled, the Incoming Relays feature is enabled globally for the appliance (relays are not listener specific). To enable the Incoming Relays feature:

1. Click the **Incoming Relays** link on the Network Tab. The Incoming Relays page is displayed.
2. Click **Enable** to enable Incoming Relays. (If the Incoming Relays feature is enabled, you can disable it by clicking Disable.)
3. Submit and commit your changes.

Adding a Relay

To add a relay:

1. Click the **Add Relay** button on the Incoming Relays page. The Add Relay page is displayed.
2. Enter a name for the relay.
3. Enter an IP Address for the relay.
4. Select a header type (Custom or Received). When entering a header, you do not need to enter the trailing colon.
5. For custom headers, enter the header name.
6. 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.
7. Commit your changes.

Filtering Your Mail Based on the Sender

There are different ways to evaluate mail as it enters your network, and one good way of determining how to filter mail is by evaluating the reputation of the person (via the IP address of the sender) who sent you mail. For example, when sorting through the envelopes that come to you via the post office, you might keep the letters sent from your grandmother, but throw away letters from *Jim's Spendalot Outlet*. You would make this decision simply by noting who sent you the mail and making judgements based on what you know about the reputation of the sender. Similarly, your Blocker can perform actions on your email based on information it can determine about the sender.

When you receive mail, the Blocker checks the IP address for each email sent to your network. You can use this IP address check to “sort” the mail that comes into your network into different groups (in much the way you might sort mail by groups of mail that come to your home: letters from people you know, letters from people you conduct business with—such as your bank, and junk mail that comes unsolicited from businesses you don't know), and you can choose to perform different actions based on these senders (IP addresses). To do this, you create a group of senders, called *Sender Groups*, and you define the behavior to perform on these senders in a *mail flow policy*. So, the combination of the sender group, plus the mail flow policy allows you to create a sorting behavior for mail that originates from certain groups of senders. For example, you might use the SenderBase reputation score to quarantine mail from senders with a questionable Senderbase reputation score.

The Host Access Table (HAT)

Each listener that is configured on your Blocker has properties that you can configure to modify the behavior of the message it receives. One of the key 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 (mail senders) 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:

Remote Host (Mail Sender) Definition	Rule
--------------------------------------	------

You can define the mail sender (remote host) in a couple of different ways--for example, by a single IP address. Once you define the mail sender, you create a rule to define whether the remote host specified can or cannot connect to the listener.

Extending the basic syntax, HATs in Blocker 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 below:

Sender Group:	Mail Flow Policy:
Sender Group:	Access Rule + Parameters
Sender 1	
Sender 2	
Sender 3	
...	

The order that rules appear in the HAT is important. The HAT is read from top to bottom for each sender that attempts to connect to the listener. If a mail sender (remote host) matches a rule, the action is taken for that sender *immediately*. This means that if the sender matches any rules further down in the HAT, those rules are ignored.

All entries you add to the HAT are entered above the final “ALL” host entry.

What is an Open Relay and Why Should I Prevent It?

For all public listeners you create, the HAT is set to accept email from all hosts by default. 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.

By rejecting all senders (other than the ones you specify), the Blocker prevents you from unintentionally configuring your system as an “open relay.” An open relay (also called an “insecure relay” or “third party relay”) is an email server that allows third-party relay of email messages. This 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 senders (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:

1. ACCEPT

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

2. REJECT

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

3. TCPREFUSE

Connection is refused at the TCP level.

4. RELAY

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

5. 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.

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.

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.
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: <code>220-hostname ESMTP</code>). You may choose to override this banner by entering a different hostname here. Additionally, you may leave the hostname field blank to choose <i>not</i> to display a hostname in the banner.
Rate Limiting	
Rate Limiting: Maximum 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.

Parameter	Description
Rate Limiting: Max. recipient per Hour Exceeded Error Code	The SMTP code returned when a host exceeds the maximum number of recipients per hour defined for this listener.
Rate Limiting: Max. Recipients Per Hour Exceeded Error Text	The SMTP banner text returned when a host exceeds the maximum number of recipients per hour defined for this listener.
Flow Control	
Use SenderBase for Flow Control	Enable “look ups” to the SenderBase Reputation Service for this listener.
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.
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 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).
Directory Harvest Attack Prevention: Drop Connection if DHAP threshold is Reached within an SMTP Conversation	The Blocker 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	
Allow TLS Connections	Deny, Prefer, or Require Transport Layer Security (TLS) in SMTP conversations for this listener.

Parameter	Description
SMTP Authentication	Allows, disallow, or requires SMTP Authentication from remote hosts connecting to the listener.
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 (RELAY only).
DKIM Verification	Enable DKIM verification.
SPF/SIDF Verification	
Enable SPF/SIDF Verification	Enable SPF/SIDF signing on this listener.
Conformance Level	Set the SPF/SIDF conformance level. You can choose from SPF, SIDF or SIDF Compatible.
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 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.
Envelope Sender DNS Verification	
	Configure Sender Verification.
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.

By default, these parameters are set to the default values shown in the table 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 Blocker 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.

Understanding the Default Sender Groups and Mail Flow Policies

Your Blocker comes with some default groups to make it easier to sort through the incoming mail. These groups are associated with mail flow policies that roughly correspond to the actions you might want to take for each type of sender.

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

Add senders to this group when you want to always reject their connections. 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 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 IronPort Anti-Spam software (if enabled for the system), the anti-virus scanning engine, and the 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.

Understanding SBRS (SenderBase Reputation Service)

The SenderBase Reputation Service (SBRS) score is a numeric value assigned to an IP address based on information from the SenderBase Reputation Service. The “reputation score” is a means of combining and averaging the information from different sources to determine the likelihood that a sender is legitimate.

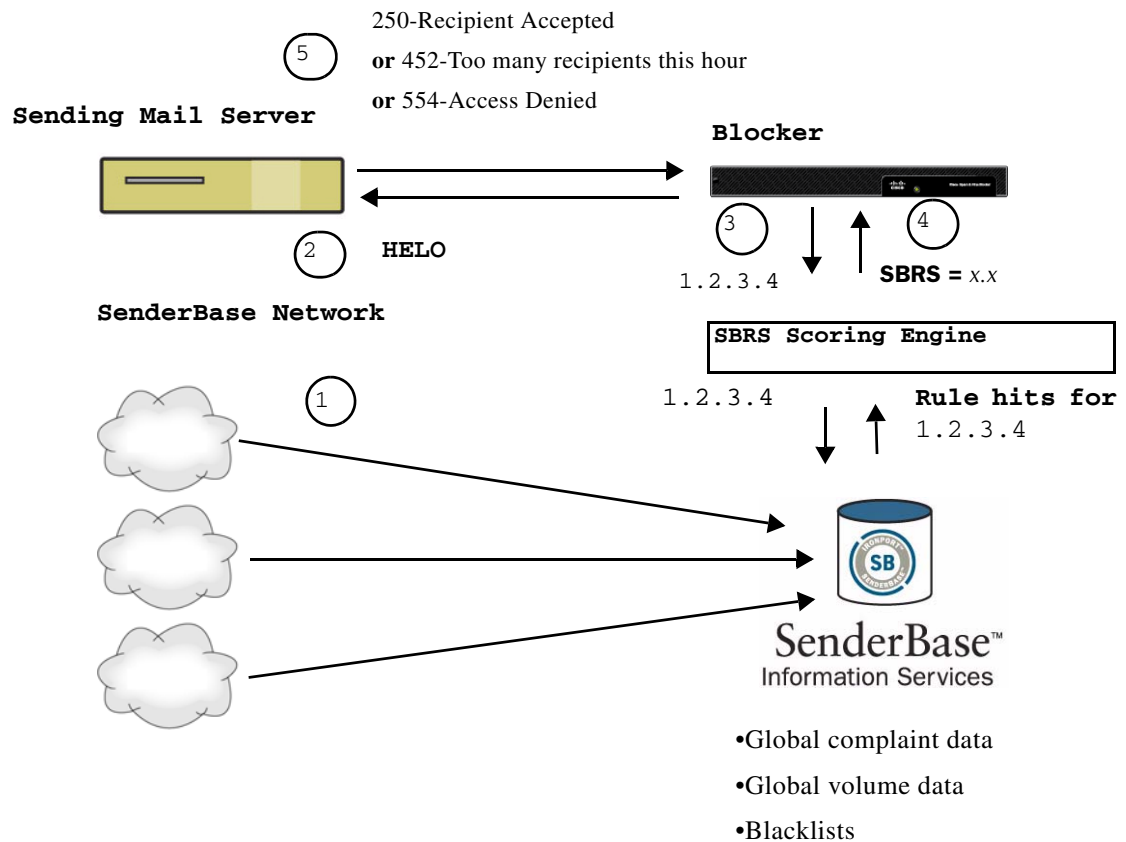
To do this, the SenderBase Reputation Service aggregates data from over 25 public blacklists and open proxy lists, and combines this data with global data from SenderBase to assign a score from -10.0 to +10.0, as follows:

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 Blocker 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.

Figure 2-1 The Senderbase Reputation Service



1. SenderBase affiliates send data
2. Sending mail server opens connection with the Blocker
3. Blocker checks global data for the connecting IP address
4. SenderBase Reputation Service calculates the probability this message is spam and assigns a SenderBase Reputation Score

- Blocker returns the response based on the SenderBase Reputation Score

Implementing SenderBase Reputation Filters

You configure reputation filtering via the **Mail Policies > HAT Overview** page. You edit these settings by entering a range of Senderbase Reputation scores from the **Edit Sender Group Settings** page.

SenderBase Reputation technology aims to shunt as much mail as possible from the remaining security services processing that is available on the Blocker. This makes filtering faster and less expensive on your system.

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%.

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.

Users are also recommended to assign all messages with a SenderBase Reputation Score greater than 6.0 to the \$TRUSTED policy.

Table 2-1 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:
Conservative:	Near zero false positives, better performance
Moderate:	Very few false positives, high performance
Aggressive:	Some false positives, maximum performance

Mail Flow Policy to Apply:
\$BLOCKED
\$THROTTLED
\$DEFAULT
\$TRUSTED

If the SBRS score returns false positives or false negatives, you can report it by following the instructions below:

- To report spam mail, forward the mail as an RFC-822 MIME encoded attachment to: *spam@access.ironport.com*.
- To report mail wrongly labeled as spam mail, forward the mail as an RFC-822 MIME encoded attachment to: *ham@access.ironport.com*.

Understanding the Email Pipeline

The Email Pipeline is the flow of email as it is processed by the Blocker. Understanding the Email pipeline is essential if you decide to configure advanced features. It is also useful for understanding how to track emails and troubleshoot problems. For example, if you are trying to find out what happened to an email sent to a mail user, understanding the email pipeline can help you to track down points at which a mail would have been rejected from the email pipeline.

In the following tables, the email pipeline is divided into two tables:

- **Acceptance.** This table represents actions that occur when mail is accepted (during the SMTP conversation).
- **Work queue.** The second table represents actions that take place in the “work queue,” after the Blocker has accepted mail and it is being processed for delivery.

It’s also important to understand when you are writing rules for processing mail. When you configure settings for the Blocker, they are always performed in a specific order.

The following table shows the actions that take place during the SMTP Conversation

Feature	Description
Host Access Table (HAT) Mail Flow Policies	<p>SMTP connections are accepted or refused based on HAT Settings:</p> <ul style="list-style-type: none"> • ACCEPT (connection is accepted) • REJECT (No email is accepted) • RELAY (Connection is accepted. Receiving for any recipient is allowed and is not constrained by the Recipient Access Table) • TCPREFUSE (Connection is refused at the TCP level.) <p>Check the queue against the TCP listen queue size (the backlog of connections that the Blocker will manage before the SMTP server accepts them)</p> <p>Perform SMTP Authentication (if enabled).</p> <p>Drop email with malformed FROM headers</p> <p>Accept or reject mail based on entries in the Sender Verification Exception Table.</p> <p>If SenderBase is enabled, Senderbase attempts to verify the IP address and reputation of the sender.</p>
Recipient Access Table (RAT)	(Public listeners only) ACCEPT or REJECT recipients in <code>RCPT TO</code> plus Custom SMTP Response. Allow special recipients to bypass throttling.
LDAP Recipient Acceptance	LDAP validation for recipient acceptance can occur within the SMTP conversation (depending on your LDAP configuration). 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.

The following actions take place in the work queue:

Feature	Description
LDAP Recipient Acceptance	LDAP validation for recipient acceptance can occur within the work queue (depending on your LDAP configuration). 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.
LDAP Masquerading	Masquerading occurs in the work queue; it rewrites the Envelope Sender, To:, From:, and/or CC: headers via an LDAP query.

Feature	Description
LDAP Routing	LDAP queries are performed for message routing or address rewriting. Group LDAP queries work in conjunction with message filter rules <code>mail-from-group</code> and <code>rcpt-to-group</code>
Safelist/Blocklist Scanning	Blocker 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* (per-recipient scanning)	Anti-Spam scanning engine examines messages and returns a verdict for further processing. * Can send messages to quarantines.
Anti-Virus* (per-recipient scanning)	Anti-Virus scanning examines messages for viruses. Messages are scanned and optionally repaired, if possible. * Can send messages to quarantines.
Content Filters* (per-recipient scanning)	Content filters are applied. You create content filters, yourself; so, there are a variety of actions that can be associated with content filters. * Can send messages to quarantines.

IMPORTANT NOTE: the above tables do not cover every possible scenario that may occur in a Blocker installation, but covers the main actions that may occur to an email as it passes through the Blocker email pipeline.

Troubleshooting Tips

This section addresses common problems that occur when running the Blocker:

- **Problem:** An email that I received should have been marked as spam. How can I get this email marked as spam?
- **Solution:** To report spam mail, forward the mail as an RFC-822 MIME encoded attachment to: *spam@access.ironport.com*.
- **Problem:** An email that I received was marked as spam when it should not have been.
- **Solution:** To report mail wrongly labeled as spam mail, forward the mail as an RFC-822 MIME encoded attachment to: *ham@access.ironport.com*.



CHAPTER 3

Mail Routing For Users and Groups

Revised: October 8, 2009 OL-20859-01

Creating Mail Rules for Groups of Users: Overview

In order to prevent your mail queue from filling with emails, you need to define the users for whom your Blocker will receive mail. However, you can also create groups of users to define different mail delivery behavior for different groups. There are many reasons you might want to create different rules for different groups of users. For example, suppose your Marketing team receives a great volume of marketing email that other teams may not want to receive. In this case, you could create a group for the marketing team that allows marketing mail and create a group for all other users that blocks or quarantines marketing mail.

There are several different ways to control mail for groups of users. First, you need to define who your listener will accept mail for. You define your users in an *LDAP directory* or using a *Recipient Access Table (RAT)*. Both of these tools tell your Blocker who you will and won't accept mail for. This helps to limit the flow of mail into your Blocker. Once you've determined who you will accept mail for, you can add further granularity by creating a *mail policy*. A mail policy allows you to create different rules for different groups of mail users. You can create very granular rules via the content filters (discussed in [Creating Custom Rules Using Content Filters, page 6-75](#)) and associate the filter with a mail policy.

Chapter Concepts/Glossary

The following concepts are discussed or introduced in this chapter. Some of these terms are specific to the Blocker appliance, while others are networking terms you will need to know in order to configure your Blocker appliance. It's a good idea to review these terms before you begin.

- **listener**. When you work with the Blocker, you configure what is called a *listener*. A listener's function is to receive messages from either internal mail servers or from external hosts. A *public* listener generally accepts connections from a sending mail server (sometimes called a *remote host*), while a *private* listener generally accepts connections from an internal mail server. To work with incoming mail for groups of users, you should have created a public listener (this is created by default when you run the System Setup Wizard).
- **Recipient Access Table (RAT)**. The Recipient Access Table defines the users whom a public listener will accept mail for. 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.

- **LDAP Server Profile.** LDAP stands for 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. When you have an LDAP directory, your Blocker can store an LDAP profile that contains connection information allowing you to run LDAP queries against an LDAP directory.
- **Mail Policies.** A mail policy is a set of rules that apply to a specific group of users and is associated with a particular listener. In this way, you can get very granular control over the rules that apply to different users. (NOTE: Mail *Flow* policies differ from mail policies in that they are designed to help you control the flow of mail on a given listener. A mail flow policy works with the Host Access Table (HAT) to define rules for mail flow).

Before You Begin

Before you create rules for your users, you should have configured listeners and created general rules for limiting the mail that flows into your Blocker via *Mail Flow Policies*, Listener configuration, SBRS configuration, and HAT configuration (covered in [Accepting Mail, page 2-29](#)). If you want to use your LDAP directory to authenticate your users, you should have set up an *LDAP User profile* and enabled LDAP on your Blocker appliance (If you have an Active Directory server, you can use the Active Directory Wizard to set up your LDAP configuration. This is covered in [Blocker Installation and Administration, page 1-1](#)). If you want to use the RAT to authenticate users, ensure that you know all the domains and other relevant information to set up your Recipient Access Table. Finally, some of the most powerful features associated with creating a group of users are the filter rules you can create for these groups. Before you create mail policies, it's a good idea to read [Creating Custom Rules Using Content Filters, page 6-75](#).

How Do I Create a List of Users?

You can define users in two places:

- **Recipient Access Table.** The Recipient Access Table defines which recipients will be accepted by the Blocker. 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.
- **LDAP Server Profile.** You can create an LDAP Server Profile to allow the Blocker to leverage the authentication settings in your LDAP Server.

By default, the Blocker creates a RAT entry with the domain you entered in the System Setup Wizard as part of your installation. You can also create an LDAP Server profile using the AD Wizard.

There are many reasons you might want to work with settings in the Recipient Access Table and LDAP authentication. The following examples shows several reasons you might want to work with these settings:

- You want to use your LDAP server to authenticate incoming mail.
- You want to accept mail for several different domains (for example, if your company acquired another company)
- You want to route mail for groups of users based on their LDAP authentication (for example, mail from members of the Marketing group as defined by the LDAP group "Marketing" might be delivered to the alternate delivery host marketingfolks.example.com).

- You want use alias expansion (LDAP routing in which Blocker replaces the original address with a new, separate email for each alias (for example, recipient@yoursite.com might be replaced with messages to newrecipient1@hotmail.com and recipient2@internal.yourcompany.com, etc.).

For example, your company, Redfish, purchases the company, Bluefish. You can configure the Blocker to receive mail for both Redfish and Bluefish domains. Similarly, you can add entries for the redfish employees in an LDAP database, and you can configure the Blocker to authenticate against this database. Finally, suppose that you want to skip LDAP authentication for Bluefish employees because the Bluefish IT group did not use LDAP authentication--you can configure the Blocker to skip LDAP authentication for this group while using LDAP authentication for the Redfish group.

Defining the Users You Accept Mail For (Using RAT)

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

Figure 3-1 Changing the RAT Entries

Order	Recipient Address	Default Action	All Delete
1	.run, ironport.com	Accept	<input type="checkbox"/>
2	redfish.com	Accept (Bypass LDAP)	<input type="checkbox"/>
	All Other Recipients	Reject	

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

For information about adding a new RAT entry, see [Adding New Mail Users or Groups of Users \(in the RAT\)](#), page 3-45

For details on how to work with the RAT, see “Configuring the Gateway to Receive Mail” in the *AsyncOS for Email User Guide*.

Adding New Mail Users or Groups of Users (in the RAT)

To add entries to a RAT:

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

Figure 3-2 Adding a Recipient to the RAT

Recipient Details	
Order:	<input type="text" value="2"/>
Recipient Address: ?	<input type="text" value="redfish.com"/>
Action:	Accept <input type="checkbox"/> Bypass LDAP Accept Queries for this Recipient <input checked="" type="checkbox"/>
Custom SMTP Response:	<input checked="" type="radio"/> No <input type="radio"/> Yes
	Response Code: <input type="text" value="250"/> Response Text: <input type="text"/>
Bypass Receiving Control: ?	<input checked="" type="radio"/> No <input type="radio"/> Yes

2. Select an order for the entry.
3. Enter the recipient address.
4. Choose to accept or reject the recipient.
5. Optionally, you can choose to bypass LDAP acceptance queries for the recipient (For details about bypassing LDAP acceptance, see “Configuring the Gateway to Receive Email” in the *AsyncOS for Email User Guide*).
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.
7. Optionally, you can choose to bypass throttling (For details about bypassing throttling for Special Recipients, see “Configuring the Gateway to Receive Email” in the *AsyncOS for Email User Guide*.) select Yes for Bypass Receiving Control.
8. Click **Submit** to add the entry.
9. Click the **Commit Changes** button, add an optional comment if necessary, then click **Commit Changes** to save the changes.

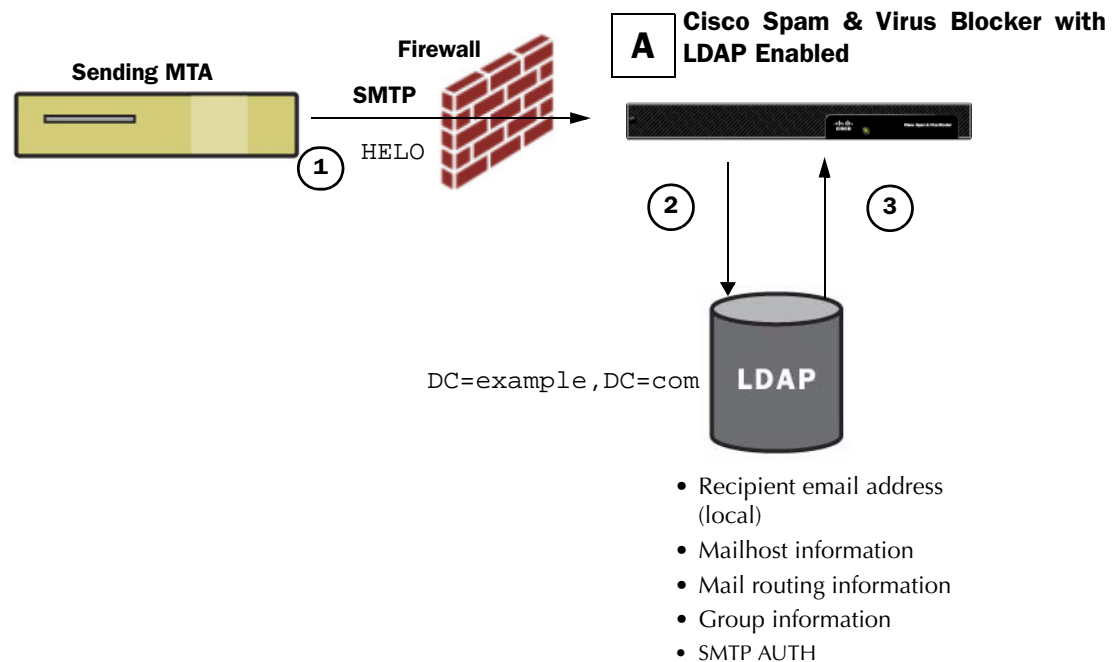
For details about deleting, importing or exporting RAT entries, see “Configuring the Gateway to Receive Email” in the *AsyncOS for Email User Guide*.

Using LDAP to Route Mail for Users and Groups of Users

When you work with LDAP directories, the Blocker can be used in conjunction with an LDAP directory server to accept recipients, route messages, and/or masquerade headers.

[Figure 3-3](#) demonstrates how the Blocker works with LDAP.

Figure 3-3 LDAP and the Blocker Appliance



1. The sending MTA sends a message to the public listener “A” via SMTP.
2. The Blocker queries the LDAP server defined via the **System Administration > LDAP** page.
3. Data is received from the LDAP directory, and, depending on the queries defined on the **System Administration > LDAP** page (or in the `ldapconfig` command) that are used by the listener:
 - the message is routed to the new recipient address, or dropped or bounced
 - the message is routed to the appropriate mailhost for the new recipient
 - From:, To:, and CC: message headers are re-written based upon the query

Configuring the Blocker to work with LDAP

When you configure your Blocker to work with an LDAP directory, you must complete the following steps to configure your Blocker appliance for acceptance, routing, aliasing, and masquerading:

1. **Configure LDAP server profiles.** The server profile contains information to enable blocker to connect to the LDAP server (or servers), such as:
 - the name of the server (s) and port to send queries,
 - the base DN, and
 - the authentication requirements for binding to the server

**Note**

If you use Microsoft Active Directory, you can use the *Active Directory Wizard* to create a server profile for your Active Directory Server. For instructions on running the Active Directory Wizard, see [Configuring Your System to Use LDAP Authentication, page 1-15](#).

For more information about configuring a server profile, see “LDAP Queries” in the *AsyncOS for Email Advanced User Guide*.

When you configure the LDAP server profile, you can configure Blocker to connect to one or multiple LDAP servers. For information about configuring Blocker to connect to multiple servers, see “LDAP Queries” in the *AsyncOS for Email Advanced User Guide*.

2. **Configure the LDAP query.** You configure the LDAP queries on the LDAP server profile. The query you configure should be tailored to your particular LDAP implementation and schema.

For information on the types of LDAP queries you can create, see “LDAP Queries” in the *AsyncOS for Email Advanced User Guide*.

For information on writing queries, see “LDAP Queries” in the *AsyncOS for Email Advanced User Guide*.

3. **Enable the LDAP server profile on a public listener or on a private listener.** You must enable the LDAP server profile on a listener to instruct the listener to run the LDAP query when accepting, routing, or sending a message.

For more information, see “LDAP Queries” in the *AsyncOS for Email Advanced User Guide*.

**Note**

When you configure a group query, you need to take additional steps to configure the Blocker to work with the LDAP server. For information on configuring a group query, see [Group LDAP Queries](#). When you configure an end-user authentication or spam notification consolidation query, you must enable LDAP end-user access to the Spam Quarantine. For more information on the Spam Quarantine, see “Configuring the Spam Quarantines Feature” in the *AsyncOS for Email User Guide*.

Understanding Mail 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 the different users within your organization.

**Note**

A powerful component of creating mail policies is the ability to create content filters that apply to specific groups. For information on creating content filters, see [Creating Custom Rules Using Content Filters: Overview, page 6-75](#).

For example, using this feature, you can quickly create policies to enforce the following conditions:

- Disable 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 drop positively identified spam. For the Human Resources organization, enable anti-spam scanning with an aggressive policy: quarantine suspected spam 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 for the word “Confidential” and words that match terms in the special code name dictionary. If a message matches, send a blind-carbon copy to the Legal department.
- 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.

**Note**

Content dictionaries, disclaimers, and notification templates must be created before they can be referenced by content filters.

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.

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.

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).

- 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, incoming messages will match different policies.

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, 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.

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, and content filters defined in policy #1, and the recipient larry@example.com will receive the anti-spam, anti-virus, 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, and content filters independently in the email pipeline.

Creating a New Policy

You can create multiple policies for different groups: for example, one for the sales organization, and another for the engineering organization. You can then configure different email security settings for each group.

1. From **Mail Policy > Incoming/Outgoing Mail Policies**, click the **Add Policy** button to begin creating a new policy.

The **Add Users** page is displayed.

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.

3. Define users for the policy.

You define whether the user is a sender or a recipient.

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 and LDAP Query

**Note**

Entries for users are case-sensitive. Use caution when entering user for a given policy. For example, if you enter the recipient Joe@ for a user, a message sent to joe@example.com will not match.

If you store user information within LDAP directories in your network infrastructure — for example, in Microsoft Active Directory, SunONE Directory Server, or Open LDAP directories — you can configure the Blocker 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.

4. 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.

5. 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.

The Mail Policies page is displayed with the new policy added.

6. Commit your changes.
7. To change anti-spam settings, click the anti-spam link.
8. To change anti-virus settings, click the anti-virus link.
9. To add content filters, click the content filters link.

**Note**

An important component of a mail policy is the associated content filter. This is discussed in [Creating Custom Rules Using Content Filters, page 6-75](#). Review this chapter to learn how to create the content filters to associate with your mail policy.

10. Submit and commit your changes.

At this point, newly created policies have the same settings applied to them as those in the default policy.



CHAPTER 4

Sending Mail

Revised: October 8, 2009, OL-20859-01

Sending Mail: Overview

To send outgoing mail from your Blocker, you need to configure a listener as a *private listener*. The private listener accepts mail from incoming mail servers such as your Exchange Email server, whereas a public listener accepts mail from outside of your network and relays it to your mail server. You must also create a *sender group* and *mail policy* for the private listener.

Use the **Network > Listeners** page to add, delete, or modify listeners. The Listeners page also provides access to the global settings for listeners.

When you create a private listener, you specify the hosts that are allowed to connect to the listener through the Host Access Table (HAT). For more information about creating HAT entries, see [Filtering Your Mail Based on the Sender, page 2-32](#).

Chapter Concepts/Glossary

The following concepts are discussed or introduced in this chapter. Some of these terms are specific to the Blocker appliance, while others are networking terms you will need to know in order to configure your Blocker appliance. It's a good idea to review these terms before you begin.

- **bounce attack.** When a message is undeliverable (typically due to a non-existent recipient address), a receiving mail system will “bounce” the email back to your mail system. Sometimes, spammers attack an email infrastructure via misdirected bounce messages. When the volume of bounced messages hits your email infrastructure, it is disruptive because email systems are not usually able to handle such a large quantity of mail.
- **bounce verification.** When enabled, IronPort Bounce Verification tags the Envelope Sender address for messages sent via your IronPort appliance. The Envelope Recipient for any bounce message received by the IronPort appliance is then checked for the presence of this tag. Legitimate bounces (which should contain this tag) are untagged and delivered. Bounce messages that do not contain the tag can be rejected.
- **private listener.** When you work with the Blocker, you configure what is called a *listener*. A listener's function is to receive messages from either internal mail servers or from external hosts. A *private* listener generally accepts connections from an internal mail server, so you'll need to create this type of listener to send mail.

- **Host Access Table.** 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.
- **mail flow policy.** 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. You'll need to create a mail flow policy for your private listener.

Before You Begin

Before you configure outbound mail, ensure that you have created a private listener to relay mail. As when you configured sender groups for incoming mail, you'll need to configure sender groups for outgoing mail. Ensure that you have the list of senders who will be authorized to send mail from your email infrastructure.

How Do I Send Mail From My Blocker?

Follow the basic steps below to configure outbound email (there are several ways to configure outbound mail, but the steps below outline the essential elements):

1. From the **Network > Listeners** page, click the HAT (Host Access Table) link for the private listener. The HAT overview page opens.
2. Ensure that your listener is selected in the Sender Groups field, and click the RELAYLIST sender group.
3. From the Edit Sender Group page, you can add senders to the RELAYLIST sender group.
4. Submit and commit your changes.

After you add your host to the relay list, the host can relay email through the Blocker.



Note

There are multiple ways to configure outbound mail. Review your Blocker documentation for additional details.

Configuring the Listener to Send Mail From Your Network

As described above, you need to create a list of users who are allowed to send mail from your network by creating entries in the **HAT** and you need to apply a **mail policy** to this group. Sender group configuration defines how a sender's IP address is "classified" (put in a sender group). Mail flow policy configuration defines how the SMTP session from that IP address is controlled. By default, if you configure a listener as private, the Blocker creates two mail flow policies for the listener: RELAYED and BLOCKED. The RELAYED policy corresponds to the sendergroup who is allowed to relay mail from the private listener, whereas the BLOCKED mail flow policy corresponds to all other mail flowing into the listener. It is important to have a policy for all other mail to prevent your Blocker from becoming an open relay.

Figure 4-1 Default Sender Group for Relay Mail Flow Policy

Sender Group Settings

Name:	RELAYLIST
Order:	1
Comment:	Only select hosts can relay from this box
Policy:	RELAYED
SBRS (Optional):	Not in use
DNS Lists (Optional):	None
Connecting Host DNS Verification:	None Included

<< Back to HAT Overview
Edit Settings...

Find Senders

Find Senders that Contain this Text: Find

Sender List: Display All Items in List

Add Sender...

There are no senders.

To relay outbound mail, senders who are allowed to relay outbound mail should be added to the RELAYLIST sendergroup.

Creating a Mail Flow Policy for Relaying Mail

By default, two mail flow policies are created for a private listener.

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

- \$RELAYED
- \$BLOCKED

Figure 4-2 Default Mail Flow Policies for Private Listeners

Sender Groups (Listener: OutgoingMail)														
Add Sender Group...		SenderBase™ Reputation Score ?					Import HAT...							
Order	Sender Group	-10	-8	-6	-4	-2	0	2	4	6	8	+10	Mail Flow Policy	Delete
1	RELAYLIST												RELAYED	🗑️
	ALL												BLOCKED	

Edit Order...
Export HAT...

By default, the following sender groups are created for a private listener:

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

Table 4-1 Default Mail Flow Policy Settings 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

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 Sender Group

Add senders you know should be allowed to relay mail 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.

Relaying Outbound Mail Via a Public Listener

You may choose to relay outbound mail through a public listener. In this case, no RelayList SenderGroup or Mail Flow Policy exists, and you will need to add them manually. To do this, select **Mail Flow Policies > Add Policy**. Assign the policy a name and choose 'Relay' from the Connection Behavior dropdown list. Submit and commit your changes.

Next, go to HAT Overview and click **Add Sender Group**, after choosing the listener from the dropdown list. Enter a name for the Sender Group, and then select the Mail Flow you recently added from the Policy dropdown list. Finally, click **Submit and Add Senders** to add your first relay host.

Adding Disclaimers to Outgoing Mail

When you send outgoing mail, you may want to include a disclaimer for mail recipients.

To add a disclaimer, you must add a text resource. All types of text resources are created in the same way, using the Text Resources page. 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 Text

The Blocker appliance can append a default text string above or below the text (heading or footer) for some or all messages received by a listener. There are three ways disclaimers can be added to messages on the Blocker appliance:

- via a listener, using the GUI. For details, see [Adding Disclaimer Text via a Listener, page 4-58](#).
- using the content filter action, Add Disclaimer Text. For details, see the [What Kinds of Actions Can I Configure?, page 6-82](#).

For example, you could append a copyright statement, promotional message, or disclaimer to every message sent from within your enterprise.

**Note**

You must have already created the disclaimer text before you can enable it. See the instructions in [To create disclaimer text:, page 4-57](#) for instructions.

To create disclaimer text:

1. go to **Mail Policies > Text Resources**.
2. Enter a name for the resource.
3. In the Type field, select **Disclaimer Text**.
4. Enter the text you want to use.
5. Add variables by selecting the **Insert Variable** link and clicking on any variables you want to use in your text.
6. Click **Submit** when you are satisfied with the message.
7. Commit your changes.

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 Blocker 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.”

Using Email Tagging to Prevent Bounce Attacks

A “bounce” message is a new message that is sent by a receiving MTA, using the Envelope Sender of the original email as the new Envelope Recipient. This bounce is sent back to the Envelope Recipient (usually) with a blank Envelope Sender (MAIL FROM: <>) when the original message is undeliverable (typically due to a non-existent recipient address).

Increasingly, spammers are attacking email infrastructure via misdirected bounce attacks. These attacks consist of a flood of bounce messages, sent by unknowing, legitimate mail servers. Basically, the process spammers use is to send email via open relays and “zombie” networks to multiple, potentially invalid addresses (Envelope Recipients) at various domains. In these messages, the Envelope Sender is forged so that the spam appears to be coming from a legitimate domain (this is known as a “Joe job”).

To combat these misdirected bounce attacks, Blocker includes Bounce Verification. When enabled, Bounce Verification tags the Envelope Sender address for messages sent via your Blocker appliance. The Envelope Recipient for any bounce message received by the Blocker appliance is then checked for the presence of this tag. Legitimate bounces (which should contain this tag) are untagged and delivered. Bounce messages that do not contain the tag can be handled separately. You can use Bounce Verification to manage incoming bounce messages based on your outgoing mail.

When sending email with bounce verification enabled, your Blocker appliance will rewrite the Envelope Sender address in the message. For example, MAIL FROM: joe@example.com becomes MAIL FROM: prvs=joe=123ABCDEFGH@example.com. The 123... string in the example is the “bounce verification tag” that gets added to the Envelope Sender as it is sent by your Blocker appliance. The tag is generated using a key defined in the Bounce Verification settings (see “Blocker Bounce Verification Address Tagging Keys” on page 83 for more information about specifying a key). If this message bounces, the Envelope Recipient address in the bounce will typically include this bounce verification tag.

You can enable or disable bounce verification tagging system-wide as a default. You can also enable or disable bounce verification tagging for specific domains. In most situations, you would enable it by default, and then list specific domains to exclude in the Destination Controls table. Use the **Mail Policies > Destination Controls** page to create, edit, and delete Destination Control entries. For more information about destination controls, see “Configuring Routing and Delivery Features” in the *AsyncOS for Email Advanced User Guide*.

Configuring Bounce Verification

To configure bounce verification, follow the steps below:

1. Add a tagging key from **Mail Policies > Bounce Verification**.
2. Enable bounce verification via Destination Controls from **Mail Policies > Destination Controls**.

3. Configure Bounce Verification from **Mail Policies > Bounce Verification**.
4. Submit and Commit your changes.

For more information details about the bounce verification settings, see *AsyncOS for Email Advanced User Guide*.



CHAPTER 5

Configuring Spam and Virus Protection

Revised: October 8, 2009, OL-20859-01

Virus and Spam Protection: Overview

Virus Scanning

The Blocker includes a virus scanning engine from Sophos, Plc. When you ran the system setup wizard, you enabled or disabled virus scanning globally and for the default mail policy you created. You can change virus scanning options, and you can also change the action that you want the Blocker to take if it discovers a virus. You can configure the Blocker to repair, archive or delete the message. You can also configure other actions-- for example, you can modify the header, add an additional descriptive header, or send a copy of the virus message to an administrator for further investigation.

Anti-spam Scanning

The Blocker includes spam scanning technology powered by IronPort Anti-Spam™ technology. When you ran the system setup wizard, you enabled or disabled anti-spam scanning settings globally for the default mail policy. However, you may want to configure anti-spam scanning for a different mail policy or even skip anti-spam scanning for certain groups. For example, you may want to skip anti-spam scanning when your partners send you mail. Like the virus-scanning engine, you can configure different actions, such as blocking a sender. You may want to send suspected spam to the IronPort Spam Quarantine, a special quarantine designed to allow your end-users to access it. This means that your end-users can create and control their own safelists and blocklists, mark certain email as “spam” and other email as “not spam,” and release messages from the quarantine to view them.

Chapter Concepts/Glossary

The following concepts are discussed or introduced in this chapter. Some of these terms are specific to the Blocker, while others are networking terms you will need to know in order to configure your Blocker appliance. It's a good idea to review these terms before you begin.

- **Mail Policy.** A mail policy is a way of grouping settings for different groups of users. You configure anti-spam and anti-virus settings in a mail policy (along with other settings such as content filters). For example, you could create a mail policy for your sales group that ignores anti-spam scanning, but maintains anti-virus settings.

- **System quarantine.** A system quarantine (unchanged from previous versions) is used to hold messages based on various actions performed by the Blocker, such as filtering and anti-virus scanning.
- **Spam quarantine.** The Spam Quarantine is a special kind of quarantine used to hold spam or suspected spam messages for end users. End users are mail users (the final recipients of the mail processed through the Blocker). Spam Quarantines can be accessed by both Blocker administrators and end users (these are not Blocker users).

Enabling Virus Protection

You should have configured global virus settings for your Blocker when you ran the system setup wizard. If you did not enable virus scanning when you ran the system setup wizard, you'll need to enable it globally before you can enable it for a mail policy. However, if you create a new mail policy, you'll need to configure Sophos Plc virus scanning settings for this policy. If enabled, virus scanning is performed in the “work queue” on the appliance, immediately after Anti-Spam scanning.

Steps to Enable Virus Protection

To manually enable and modify the virus scanning engine global configuration settings, go to **Security Services > Sophos Anti-Virus** pages (GUI).

To enable Anti-Virus scanning if you have not previously enabled an anti-virus engine in the System Setup Wizard, complete the following steps:

1. Select **Security Services > Sophos Anti-Virus**
2. Click **Enable**. The license agreement page is displayed.
Clicking **Enable** enables the feature globally for the appliance. However, you must later enable per-recipient settings in Mail Policies.
3. After reading the agreement, scroll to the bottom of the page and click **Accept** to accept the agreement.
4. Click **Edit Global Settings**.
5. Choose a maximum virus scanning timeout value.
6. Configure a timeout value for the system to stop performing anti-virus scanning on a message. The default value is 60 seconds.
7. Click **Submit**. The **Security Services > Sophos Anti-Virus** page is refreshed to display the values you chose in the previous steps.

Editing the Anti-Virus Settings for a Mail Policy

The process for editing the 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 **Mail Policies > Incoming or Outgoing Mail Policies** pages. 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.

To edit the anti-virus settings for a mail policy, including the default policy:

1. Click the link for the anti-virus security service in any row of the Email Security Manager incoming or outgoing mail policy table.
2. Click the link in the default row to edit the settings for the default policy.
3. 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.

4. Enable the Sophos Anti-virus scanning engine.
5. Configure Message Scanning settings.
6. Configure settings for Repaired, Encrypted, Unscannable, and Virus Infected messages.
7. Configure Advanced settings for Repaired, Encrypted and Virus infected messages if you want to create a custom header, or create a custom notification for users when a message is tagged as repaired, encrypted, or virus positive. For instructions on creating a custom notification, see [Anti-Virus Notification Templates, page 5-63](#).



Note You must have already created the notification template to select it from the mail policy. If you want to use custom notifications, create them before you configure your mail policy.

8. Click **Submit**.

The **Mail Policies > Incoming or Outgoing Mail Policies** page is refreshed to reflect the values you chose in the previous steps.

Click the **Commit Changes** button, add an optional comment if necessary, and then click **Commit Changes** to finish configuring Anti-Virus settings for a Mail Policy.



Note

When the Blocker detects a virus, it is often configured to send the virus message to a system quarantine. If you send a message to a system quarantine, see the instructions below for accessing the quarantine and releasing messages.

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.

You can specify a custom notification to send while editing a mail policy. You can configure the From: address for anti-virus notifications.

To create an anti-virus notification:

1. go to **Mail Policies > Text Resources**.
2. Enter a name for the resource.
3. Select a type: you can choose one of the following:
 - Anti-virus container template
 - Anti-virus notification template
4. Enter the text you want to use.
5. Add variables by selecting the **Insert Variable** link and clicking on the variables you want to use in your text.
6. Click submit when you are satisfied with the message.
7. Commit your changes.

Working with System Quarantines

A system quarantine is used to hold messages based on various actions performed by the Blocker, such as filtering and anti-virus scanning.

Typically, messages are placed in system quarantines due to a filter action. System quarantines are configured to process messages automatically—messages are either delivered or deleted based on the configuration settings set for the quarantine(s) in which the message is placed. In addition to the automated process, designated users (such as your mail administrator, Human Resources personnel, Legal department, etc.) can review the contents of the quarantines and then either release, delete, or send a copy of each message. Released messages are scanned for viruses (assuming that anti-virus is enabled for that particular mail policy).

Your Blocker can have several pre-configured quarantines, depending on features licensed; however, the Policy quarantine is created by default, regardless of license.

- Virus, a quarantine used by the anti-virus engine, created when the anti-virus license key is enabled.
- Policy, a default quarantine (for example, use this to store messages requiring review).

System Quarantine Settings

Quarantines have an automated process for handling messages based on quarantine settings. Quarantines have several settings used to determine how the quarantine acts on a day-to-day basis: Space Allocation, Retention Time, Default Action, Overflow Messages, and Users. Once you have made a change, submit and commit your changes.

Retention Time

Retention Time is the length of time messages are kept in a quarantine. The Default Action (see Default Action) is performed on any message in the quarantine once that retention time is reached. Each message has its own specific expiration time, displayed in the quarantine listing.

Messages are stored for the amount of time specified (Normal Expiration) unless they are manually processed by a mail administrator (or other user) or the size limit set for the quarantine is reached. If the size limit is reached, the oldest messages are processed (Early Expiration) and the Default Action is performed for each message until the size of the quarantine is again less than the size limit. The policy is First In First Out (FIFO).

The expiration time on a message can be delayed (extended) via the **Select Action** menu in the various quarantine listings. Delaying the expiration of a message can be helpful when you need to keep specific messages in the quarantine past their scheduled expiration (for example, waiting for an administrator to have time to review the messages or for a specific anti-virus IDE to be published).

Default Action

The Default Action is the action performed on messages in a quarantine when either of the two following circumstances occur:

- **Normal Expiration** - the Retention Time is met for a message in the quarantine (see Retention Time).
- **Early Expiration** - a message is forced from the quarantine when the size limit for the quarantine is reached. For more on setting size limits for quarantines, see *Creating System Quarantines*. Messages released from quarantine because of a queue-full condition (early expiration) can optionally have other operations performed on them. For more information, see *Overflow Messages*.

There are two Default Actions:

- **Delete** - the message is deleted.
- **Release** - the message is released for delivery. Upon release, the message is rescanned for viruses, assuming anti-virus is enabled for that particular mail policy. For more information about virus scanning and messages released from quarantines, see *System Quarantines and Virus Scanning*.



Note

In addition to these two default actions, a third message action (Delay Exit) is available in the Select Action menu in the quarantined messages listing.

Overflow Messages

The Overflow Messages section is used to dictate how messages are handled as they are released from the quarantine due to overflow. These settings include: Subject Tagging, Adding an X-Header, and Stripping Attachments.

Subject Tagging

Messages released or deleted from a quarantine because of a queue-full condition (early expiration only) can optionally have their subjects tagged with text you specify when editing or creating a quarantine.

The tag is a user-defined string that can either be prepended or appended to the original subject header.

Note — In order for a subject with non-ASCII characters to display correctly it has to be represented according to RFC 2047.

Add X-Header

Messages released or deleted from a quarantine because of a queue-full condition (early expiration only) can optionally have an X-Header added.

Specify the name of the X-Header and the value.

Strip Attachment

Messages released or deleted from a quarantine because of a queue-full condition (early expiration only) can optionally have their attachments stripped. This can be used to help reduce the chance for virus infected files will be released from a quarantine.

Users and User Groups

Users belonging to the Administrators group have access to quarantines by default. Users in the Operators, Guests, Read-Only Operators, and Help Desk Users groups can be assigned to a quarantine (so that they may view, process, or search messages within a quarantine), but cannot change the quarantine's configuration (e.g. the size, retention period, etc.), or create or delete quarantines.

Creating System Quarantines

You can create new system quarantines to hold messages. The basic workflow for setting up a quarantine is:

- Create local users that will interact with the quarantine. A quarantine's user list contains local users in all user groups, except Administrators. Users in the Administrators group always have full access to the quarantine. For more information, see [Adding Users](#). You can also enable your IronPort appliance to use an external directory to authenticate users and select which user groups have access to the quarantine. For more information, see [External Authentication](#).
- Create the quarantine, following the steps below.
- Create filters that will move messages to the quarantine. For more information about creating filters, see [Creating Custom Rules Using Content Filters: Overview, page 6-75](#).

To create a system quarantine:

1. Click **Add Quarantine** on the Quarantines page. The Add Quarantine page is displayed.
2. Type a name for the quarantine.
3. Specify the space (in megabytes) to allocate for the quarantine. For more information, see [Allocating Space for System Quarantines](#).
4. Select a Retention Period, or time to keep a message in the quarantine before the default action is performed on the message. For more information, see [Retention Time](#).
5. Select a Default Action (Delete or Release).
6. If you want to modify the subject of messages that are processed through the quarantine, type the text to add and select whether to prepend or append it to the original message subject. For more information, see [Subject Tagging](#).
7. If you want to add an X-Header, enter the name and value. For more information, see [Add X-Header](#).

8. If you want to strip any file attachments when the file is release from the quarantine due to overflow (early release), select On. For more information, see Strip Attachment.
9. Select users to associate with this quarantine by clicking on the name of the user. Hold the CTRL key and click to select multiple users. A quarantine's user list contains local users in all user groups, except Administrators. Users in the Administrators group always have full access to the quarantine. For more information, see Users and User Groups. A warning is displayed if you have not yet created any users.
10. Optionally, select the check boxes for the user roles of externally authenticated users to associate with the quarantine. Externally authenticated users are authenticated by your IronPort appliance using a centralized authentication system. For more information, see External Authentication.
11. Submit and commit your changes.

Editing System Quarantines

Only users belonging to the Administrators group can edit quarantines.

To edit an existing quarantine:

1. Click the **Edit** link in the Settings column for the quarantine you want to modify. The Edit Quarantine page is displayed:
2. Make your changes to the settings for the quarantine.
3. Sumit and commit your changes.

Deleting System Quarantines

To delete an existing quarantine:

1. Click the Delete Quarantine link in the Edit Quarantine page.
2. A confirmation message is displayed:
3. Click **Delete**. The quarantine is deleted.
4. Commit your changes.

Enabling Spam Protection

The Blocker provides a layered approach to stopping spam. 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 IronPort SenderBase™ Reputation Service. The second layer of defense, scanning, is powered by 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 Blocker, 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.

Steps to Enable Anti-Spam Scanning

To enable Anti-Spam, follow these steps:

1. If you have not enabled Anti-Spam in the system setup wizard, select **Security Services > IronPort Anti-Spam**.
2. Click **Enable**.

The license agreement page is displayed.

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

3. Scroll to the bottom of the page and click **Accept** to accept the agreement. Click **Edit Global Settings**.
4. 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.
5. Choose a value for the maximum message size to scan by IronPort Anti-Spam.
The default value is 128 Kb. Messages larger than this size will not be scanned by IronPort Anti-Spam and the *X-IronPort-Anti-Spam-Filtered: true* header will not be added to the message.
6. Enter a length of time to wait for timeout when scanning a message.
When entering a time value, use 's' for seconds, 'm' for minutes, and 'h' for hours. For example '1m' is one minute. The default value is 1 minute.
7. Enable or disable regional scanning. Regional scanning optimizes 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.
8. Submit and commit your changes.

The **Security Services > IronPort Anti-Spam** page is refreshed to display the values you chose in the previous steps.

Configuring the Spam Quarantine

The Spam Quarantine is a special kind of quarantine used to hold spam or suspected spam messages for end users. End users are mail users (the final recipients of the mail processed through the Blocker). You can have a local Spam Quarantine, stored on the Blocker. Spam Quarantines can be accessed by both Blocker administrators and end users (these are not Blocker users).

Each Blocker can have a local Spam Quarantine enabled if the IronPort Anti-spam scanning has been enabled. Follow these steps to configure your Blocker to send spam or suspect spam messages to a Spam Quarantine:

- **Enable and configure the local Spam Quarantine.** Configuring the local Spam Quarantine allows you to specify settings related to quarantine access, contents, and behavior, notifications, authentication, and Blocker user access.
- **Edit an IP interface and enable the Spam Quarantine HTTP or HTTPS service.** Enabling the Spam Quarantine HTTP/S service allows you to access the quarantine.

Enabling and Disabling the Spam Quarantine

Follow the steps below to enable the Spam Quarantine.

1. On the **Monitor > Quarantines** page, click **Enable**.

Figure 5-1 Spam Quarantine Main Page
Quarantines

Quarantines				
Add Quarantine...		No space available for additional quarantines.		
Quarantine	Messages	Default Action	Status	Settings
Spam Quarantine (disabled) **	--	--	--	Enable
Policy	0	Retain 10 days then Delete	<input type="text"/> 0% Full	Edit
Virus *	0	Retain 30 days then Delete	<input type="text"/> 0% Full	Edit

* This Quarantine cannot be used until the related security service is enabled.
** This Quarantine cannot be used until the Spam Quarantine HTTP or HTTPS service is enabled on one of your IP Interfaces. Go to Network > IP Interfaces to configure this.

2. The Spam Quarantine is enabled. If the Spam Quarantine is not configured, the Edit Spam Quarantine page is displayed.
3. Submit and commit your changes.

Disabling the Spam Quarantine

To disable the local Spam Quarantine:

1. On the **Monitor > Quarantines** page, click **Edit** in the Settings column for the Spam Quarantine.
2. In the Spam Quarantine Settings section, uncheck **Enable Spam Quarantine**.
3. Submit and commit your changes.

If messages are present in the local Spam Quarantine when it is disabled, you can opt to delete all of the messages via the “Delete All” link on the Quarantines page.

Configuring Settings for the Spam Quarantine

To edit the Spam Quarantine settings:

1. Click **Edit** in the Settings column for the Spam Quarantine on the **Monitor > Quarantines** page. The Edit Spam Quarantine page is displayed.
2. In the Spam Quarantine Settings section, specify a maximum quarantine size.
You can configure the quarantine to delete the oldest messages when the quarantine is full. If unchecked, newer messages will not be added to a full quarantine. Cisco recommends that you enable this feature so that a full quarantine will not cause messages to queue (back up) on your appliance.
3. Specify the number of days to hold messages before deleting them, or you can elect to not schedule automatic deletion. Cisco recommends that you configure the quarantine to delete older messages to prevent the quarantine from filling to capacity.
4. Specify a default language.

You can configure the quarantine to send a copy of released messages to Cisco for analysis. Cisco recommends that you do configure the quarantine to do so.

5. Customize the page end users see when they view the quarantine. Upload a custom logo (optional). The logo is displayed at the top of the Spam Quarantine page when the user logs in to view quarantined messages.
 - The logo should be a .jpg, .gif, or .png file that is at most 550 x 160 pixels.
 - If a logo file is not supplied, the default Spam Quarantine logo is used.



Note If you specify a custom logo, the Cisco logo is deleted.

6. Specify a login page message. This message is shown to end users when they are asked to log in prior to viewing the quarantine.
7. Submit and commit your changes.

Enabling End User Safelists and Blocklists

You can enable end users to create safelists and blocklists to better control which emails are treated as spam. Safelists allow a user to ensure that certain users or domains are never treated as spam, while blocklists ensure that certain users or domains are always treated as spam. The safelists and blocklists settings are configured from the Spam Quarantine, so you must enable and configure the Spam Quarantine to use this feature. When you enable the safelist/blocklist feature, each end user can maintain a safelist and blocklist for his or her email account.

Safelists and blocklists prevent mail from being treated as spam or ensure that mail is always treated as spam. However, a safelist or blocklist setting does not prevent the Blocker from scanning an email for viruses or determining if the message meets the criteria for a content-related mail policy. If a message is part of a safelist, it may not be delivered to the end user depending on other scanning settings. For example, if a user configures a mail sender on his safelist, but the sender's mail contains a virus, then the mail may be quarantined or blocked.

Creating and Maintaining Safelists and Blocklists

The safelists and blocklists are created and maintained by end users. However, an administrator enables the feature and configures delivery settings for email messages matching entries in the blocklist. To create and maintain safelists and blocklists, the administrators and end-users complete the following tasks:

- **Administrator tasks.** Administrators enable and configure the Spam Quarantine, enable the Safelist/Blocklist feature, backup and restore the Safelist/Blocklist database, synchronize the Safelist/Blocklist database between different appliances, and troubleshoot safelist and blocklist issues via logs, alerts, and custom headers.
- **End-user tasks.** End-users create their safelist and blocklist settings via the end-user spam quarantine. End users may need to log in (instead of clicking the link in the Spam Quarantine notification) to access their safelist/blocklist settings. From the end-user spam quarantine, end-users can create safelists and blocklists from the Options menu. Or, end-users can create safelist settings from the list of quarantined emails.

Administrator Tasks for Creating and Maintaining Safelists and Blocklists

To use safelists and blocklists, the administrator must complete the following tasks:

- **Enable and configure the Spam Quarantine.** Because the safelist and blocklist is accessed from the Spam Quarantine, you must enable this feature to use safelists and blocklists.
- **Enable and configure the Safelist/Blocklist feature.** Once the Spam Quarantine is enabled, you enable and configure the Safelist/Blocklist feature. You must also configure a blocklist action for blocklisted email (quarantine or delete).
- **Backup and restore the Safelist/Blocklist database.** When upgrading, you need to backup and restore the Safelist/Blocklist database.
- **Troubleshooting Safelists and Blocklists.** To troubleshoot safelists and blocklists, you can check logs and alerts.

Enabling and Configuring Safelist/Blocklist Settings

You can enable and configure settings for safelists and blocklists from the Quarantines page.

1. To enable safelists and blocklists on a Blocker appliance, go to **Monitor > Quarantines**.
2. In the End-User Safelist/Blocklist settings, select Edit Settings.



Note

You must have the Spam Quarantine enabled and configured before you can configure safelists and blocklists.

3. Select Enable Safelist/Blocklist Feature.
4. Select Quarantine or Delete for the Blocklist Action.
5. Specify the Maximum List Items Per User. This value represents the maximum number of addresses or domains a user can list in each safe and block list.
6. Click **Submit**.

Backing Up and Restoring the Safelist/Blocklist Database

To save a backup of the safelist/blocklist database, the Blocker saves the database as a .CSV file. The .CSV file is maintained separately from the XML configuration file that contains your Blocker configuration settings. If you upgrade your Blocker or run the system setup wizard, you should back up the Safelist/Blocklist database to the .CSV file.

When you back up a file, the Blocker saves a .CSV file to the /configuration directory using the following naming convention:

```
slbl<timestamp><serial number>.csv
```

From the GUI, you can use the following method to back up and restore the database:

1. From **System Administration > Configuration File**, go to the End-User Safelist/Blocklist Database section.
2. To back up a database to a .CSV files, click **Backup Now**.
3. To restore the database, click **Select File to Restore**.
4. The Blocker displays a list of backup files that are stored in your configuration directory.

5. Select the safelist/Blocklist backup file you want to restore and click **Restore**.

Troubleshooting Safelists and Blocklists

An end user maintains his or her own safelists and blocklists. Administrators can access an end user's safelist or blocklist only by logging into the end user account with the user's login and password. To troubleshoot issues with safelists and blocklists, you can view the log files or system alerts.

When an email is blocked due to safelist/blocklist settings, the action is logged in the EUQ_logs or the antispam log files. Emails that are safelisted are marked as safelisted with an *X-SLBL-Result-Safelist* header. Emails that are blocklisted are marked as blocklisted with an *X-SLBL-Result-Blocklist* header.

Alerts are sent out when the database is created, updated, or if there are errors in modifying the database or running the safelist/blocklist processes.

For more information about alerts, see “Alerts” in the *AsyncOS for Email User Guide*.

For more information about log files, see “Logging” in the *AsyncOS for Email Advanced User Guide*.

End User Tasks for Configuring Safelists and Blocklists

End users can create safelists to ensure that messages from certain senders are never treated as spam, and they can use blocklists to ensure that messages from certain senders are always treated as spam. For example, an end user may receive email from a mailing list that no longer interests him. He may decide to add this sender to his blocklist to prevent emails from the mailing list from getting sent to his inbox. On the other hand, end users may find that emails from specific senders get sent to their Spam Quarantine when they don't want them to be treated as spam. To ensure mail from these senders are not quarantined, they may want to add the senders to their safelists.



Note

Safelist/Blocklist settings are contingent on other settings configured by the system administrator. A safelisted message may not be delivered if it is determined to be virus-positive, or if the administrator determines that the content does not conform to company email policies.

To work with safelists and blocklists, end users must complete the following tasks:

- **Access safelists and blocklists.** Depending on authentication settings, end users may need to log into their Spam Quarantine accounts.
- **Add safelist entries.** Users add safelist entries from the Options menu or the list of quarantined messages in Spam Quarantine.
- **Add blocklist entries.** Users add blocklist entries from the Options menu of the Spam Quarantine.

Accessing Safelists and Blocklists

To access safelists and blocklists, end users whose accounts are authenticated using LDAP or Mailbox (IMAP/POP) authentication must log into their accounts on the Spam Quarantine. The end user must log into their account even if they are accustomed to accessing their messages via a spam notification (which usually doesn't require authentication). If the end-user authentication is set to NONE, end users do not need to log into their accounts to access safelist/blocklist settings.

Syntax for Safelists and Blocklist Entries

Entries can be added to safelists and blocklists using the following formats:

user@domain.com

server.domain.com

domain.com

End users cannot add a sender or domain to both safe and block lists at the same time. However, if the end user adds a domain to a safelist, and the email address for a user of that domain to the blocklist (or vice versa), the Blocker applies both rules. For example, if the end user adds *example.com* to the safelist, and adds *george@example.com* to the blocklist, the Blocker delivers all mail from *example.com* without scanning for spam, but will treat mail from *george@example.com* as spam.

End users cannot allow or block a range of sub-domains using the following syntax: *.domain.com*. However, an end user can explicitly block a specific domain using the following syntax: *server.domain.com*.

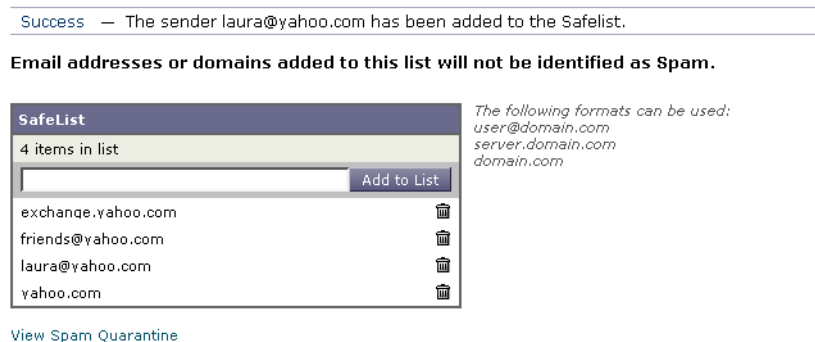
Adding Entries to Safelists

End users can add senders to safelists in two ways:

Method 1:

1. From the Spam Quarantine, select the Options drop-down menu.
2. Choose Safelist.
3. From the Safelist dialog box, enter the email address or domain.
4. Click **Add to List**.

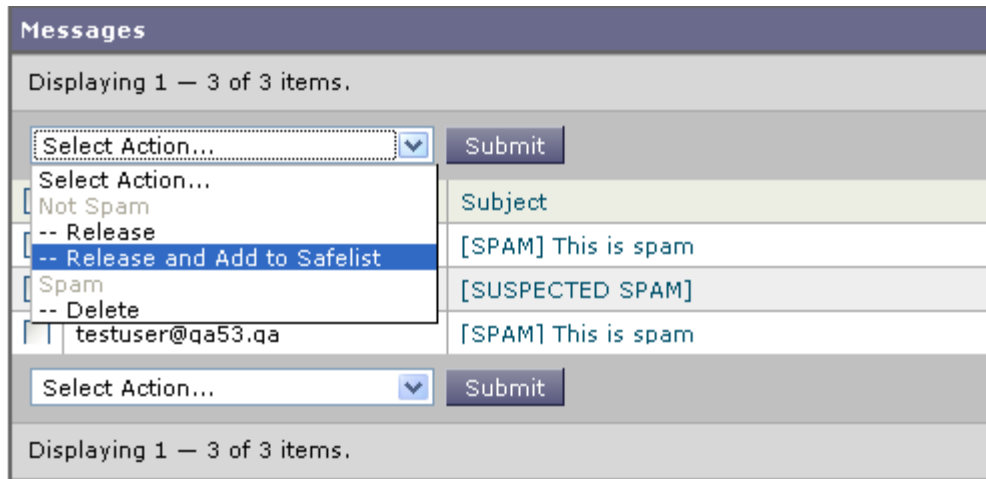
Figure 5-2 Release and Add to Safelist



Method 2:

1. End users can also add senders to the safelist if the message has been sent to the end user quarantine.
2. From the End-User Quarantine, select the checkbox next to message.
3. Choose "Release and Add to Safelist" from the drop-down menu.

Figure 5-3 Safelist in End User Quarantine



The envelope sender and the from header for the specified mail are both added to the safelist, and the released messages proceed directly to the destination queue, skipping any further work queue processing in the email pipeline.

Adding Entries to Blocklists

1. End users can use blocklists to ensure that they never receive mail from specified senders.
2. From the End-User Quarantine, select the Options drop-down menu.
3. Enter the domain or email address you want to blocklist.
4. Click **Add to List**.

When the Blocker receives mail from the specified email address or domain that matches an entry in the blocklist, it treats the mail as spam. The mail may be rejected or it may be quarantined, depending on the safelist/Blocklist action settings.



Note

Unlike safelist entries, you can only add blocklist entries from the Options menu in the End-User Quarantine.



CHAPTER 6

Creating Custom Rules Using Content Filters

Revised: October 8, 2009, OL-20859-01

Creating Custom Rules Using Content Filters: Overview

Just as you create custom rules for your email inbox (such as creating filters to send all emails from the engineering department to a special folder), you will likely want to create some special rules for your Blocker appliance. One way to do this is to create a content filter. For example, you might want to search for terms that violate your company's privacy policies. You can flag incoming mail that contain key words (or key patterns--such as the numeric pattern of a social security number) and send a message to the Email Administrator at the same time. Or, you may want to ensure that you add a disclaimer to sensitive outgoing mail. You can create a filter that searches for certain parameters (such as a particular sender or group), and add a disclaimer to all messages sent from these mail users. To enforce these rules you create a **content filter**.

Chapter Concepts/Glossary

The following concepts are discussed or introduced in this chapter. Some of these terms are specific to the Blocker appliance, while others are networking terms you will need to know in order to configure your Blocker appliance. It's a good idea to review these terms before you begin.

- **Multipurpose Internet Mail Extensions (MIME)**. MIME refers to an official Internet standard that specifies how messages must be formatted so that they can be exchanged between different email systems. When you apply content filters, you can apply them to different parts of a message (when the message is MIME encoded). Ensure you know basic MIME format to search for files, file extensions or terms in the right area. Also, be aware that for the purpose of content filters, Cisco treats anything after the message body as an attachment, even though this differs from the MIME specification. This is consistent with how many email programs present messages and attachments.
- **Mail Policies**. A mail policy is a set of rules that apply to a specific group of users and is associated with a particular listener. To use content filters you must enable them on a mail policy (which will apply to a group of users on a particular listener). In this way, you can get very granular control over the rules that apply to different users. (NOTE: Mail *Flow* policies differ from mail policies).
- **smart identifiers**. When you use message rules that scan message content, you can use smart identifiers to detect certain patterns in the data. Smart identifiers can detect the following patterns in data:
 - Credit card numbers

- U.S. Social Security numbers
- Committee on Uniform Security Identification Procedures (CUSIP) numbers
- American Banking Association (ABA) routing numbers
- **Threshold scoring.** When you add content 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 the Blocker 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 content filter does not perform the configured action on the message. The message scoring works slightly differently depending on whether the message is MIME encoded or not. Threshold scoring also applies to content dictionaries.

Before You Begin

You apply content filters by enabling them on a *mail policy*. You must have already created a mail policy in order to enable the content filters, so ensure that you have configured at least one mail policy and that you understand how they work before you enable content filters. For information about creating and implementing mail policies, see [Policy Matching, page 3-49](#).

You may want to send some messages to the *system quarantine*. Ensure that you understand how quarantines work if you intend to route mail to a system quarantine.

When you perform searches for different terms, the Blocker uses *threshold scoring* to allow you to give different weights to different terms. This is a way of giving you some control over the importance of different terms. But in order to effectively use threshold scoring, it’s a good idea that you understand how scoring works with the different MIME (Multipurpose Internet Mail Extensions) message parts. For more information, see [Understanding Threshold Scoring in Content Filtering, page 6-88](#).

In addition, some of the content filters use a *content dictionary*, which defines terms to search for (for example, you could create a dictionary called “banking terms” that includes terms related to banking. You must create the content dictionary *before* you can use it in the content filters. For instructions on creating a content dictionary, see [How Do I Create a List of Terms to Search For?, page 6-89](#).

Understanding How Content Filters Work

You create content filters to be applied to messages on a per-recipient or per-sender basis. Content filters are applied 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 anti-spam and anti-virus scanning have been performed on a message. This means that some messages may have been filtered or bounced before the content filter can be applied to it.

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 in the order in which they will be applied for any given policy.

How Do I Create a Content Filter?

The basic steps for creating a content filter are as follows:

1. Click the **Mail Policies** tab.
2. Click the **Incoming Content Filters** Section.

The Incoming Content Filters page is displayed. On newly installed or upgraded systems, no content filters are defined by default.

3. Click the **Add Filter** button.

The Add Content Filter page is displayed.

4. Enter a name and description for the filter.

5. Click **Add Condition**.

For information about the conditions you can apply, see “How Do I Add Conditions to My Filter?” on page 77.

6. Select a condition to add.

7. Click **Add Action**.

For information about the actions you can apply, see “What Kinds of Actions Can I Configure?” on page 82.

8. Click **OK**.

At this point, the content filter is not enabled for any incoming Mail Policy. Because it has not been applied to any policy, no email processing will be until you apply it to a Mail Policy. For instructions about creating and enabling a mail policy on a listener, see [Understanding Mail Policies, page 3-48](#).

How Do I Add Conditions to My Filter?

Specifying conditions in content filters is optional. If you don't specify a condition, then the filter will be applied to all mail policies associated with the content filter.

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 Blocker 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 *AsyncOS for Email Advanced User Guide*.

Multiple conditions may be 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”).

An email message is composed of multiple parts (header, body, attachment, envelope sender, envelope recipient). Although Rafts define everything that comes after a message’s headers as a multipart “message body,” many users still conceptualize a message’s “body” and its “attachment” differently. For the purposes of writing body or attachment content filter rules, everything after the message headers is considered the message body, whose content is considered the first text part of the MIME parts that are within the body. Anything after the content, (that is, any additional MIME parts) is considered an attachment.

The following shows the content filter condition screen:

Figure 6-1 Content Filter Conditions

Add Condition

Message Body or Attachment

Message Body
Message Size
Attachment Content
Attachment File Info
Attachment Protection
Subject Header
Other Header
Envelope Sender
Envelope Recipient
Receiving Listener
Remote IP
Reputation Score
DKIM Authentication
SPF Verification

Message Body or Attachment [Help](#)

Does the message body or attachment contain text that matches a specified pattern?

Contains text:
*

Contains smart identifier:

Contains term in content dictionary:
No content dictionaries are defined. See Mail Policies > Dictionaries.

Number of matches required: (1-1000)
For content dictionaries, the number of matches is based on term weight.

(*) accepts regular expression

Table 6-1 Content Filter Conditions

Condition	Description
(no conditions)	Specifying conditions in content filters is optional. If no conditions are specified, a <code>true</code> rule is implied. The <code>true</code> rule matches all messages, and the actions are always performed.
Message Body or Attachments	<p>Contains text: Does the message body contain text or an attachment that matches a specific pattern?</p> <p>Contains smart identifier: Does content in the message body or attachment match a smart identifier?</p> <p>Contains term in content dictionary: Does the message body contain any of the regular expressions or terms in the content dictionary named <i><dictionary name></i>?</p> <p>For this option to be enabled, the dictionary must already have been created.</p> <p>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.</p> <p>This includes delivery-status parts and associated attachments.</p>
Message Body	<p>Contains text: Does the message body contain text that matches a specific pattern?</p> <p>Contains smart identifier: Does content in the message body match a smart identifier?</p> <p>Contains term in content dictionary: Does the message body contain any of the regular expressions or terms in the content dictionary named <i><dictionary name></i>?</p> <p>For this option to be enabled, the dictionary must already have been created.</p> <p>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.</p> <p>This rule applies to the body of the message only. It does not include attachments or headers.</p>

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	<p>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 <code>body-contains()</code> 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.</p> <p>Contains a smart identifier. Does content in the message attachment match the specified smart identifier?</p> <p>Contains terms in content dictionary. Does the attachment contain any of the regular expressions or terms in the content dictionary named <i><dictionary name></i>?</p> <p>To search for dictionary terms, the dictionary must already have been created.</p> <p>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.</p>
Attachment File Info	<p>Filename. Does the message contain an attachment with a filename that matches a specific pattern?</p> <p>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 <code>file</code> command)?</p> <p>MIME type. Does the message contain an attachment of a specific MIME type? This rule is similar to the <code>attachment-type</code> 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.)</p>

Condition	Description
Attachment Protection	<p>Contains an attachment that is password-protected or encrypted.</p> <p>(for example, use this condition to identify attachments that are potentially unscannable)</p> <p>Contains an attachment that is NOT password-protected or encrypted.</p> <p>(For example, use this condition with the Encrypt action to make sure all attachments are encrypted.)</p>
Subject Header	<p>Subject Header: Does the subject header match a certain pattern?</p> <p>Contains terms in content dictionary: Does the subject header contain any of the regular expressions or terms in the content dictionary <i><dictionary name></i>?</p> <p>To search for dictionary terms, the dictionary must already have been created.</p>
Other Header	<p>Header name: Does the message contain a specific header?</p> <p>Header value: Does the value of that header match a certain pattern?</p> <p>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 <i><dictionary name></i>?</p> <p>To search for dictionary terms, the dictionary must already have been created.</p>
Envelope Sender	<p>Envelope Sender. Does the Envelope Sender (i.e., the Envelope From, <MAIL FROM>) match a given pattern?</p> <p>Matches LDAP group. Is the Envelope Sender, i.e., the Envelope From, <MAIL FROM> in a given LDAP group?</p> <p>Contains term in content dictionary. Does the envelope sender contain any of the regular expressions or terms in the content dictionary named <i><dictionary name></i>?</p> <p>To search for dictionary terms, the dictionary must already have been created.</p>

Condition	Description
Envelope Recipient	<p>Envelope Recipient. Does the Envelope Recipient, (i.e. the Envelope To, <RCPT TO>) match a given pattern?</p> <p>Matches LDAP group. Is the Envelope Recipient, (i.e. the Envelope To, <RCPT TO>) in a given LDAP group?</p> <p>Contains term in content dictionary. Does the envelope recipient contain any of the regular expressions or terms in the content dictionary named <<i>dictionary name</i>>?</p> <p>To search for dictionary terms, the dictionary must already have been created.</p> <p>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.</p> <p>Is the Envelope Sender (i.e., the Envelope From, <MAIL FROM>) in a given LDAP group?</p>
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.
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, see <i>AsyncOS for Email Advanced User Guide</i> .

What Kinds of Actions Can I Configure?

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 the *AsyncOS for Email Advanced User 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 will force final actions to be placed last.

The following shows the Content Filter Actions screen:

Figure 6-2 Content Filter Actions

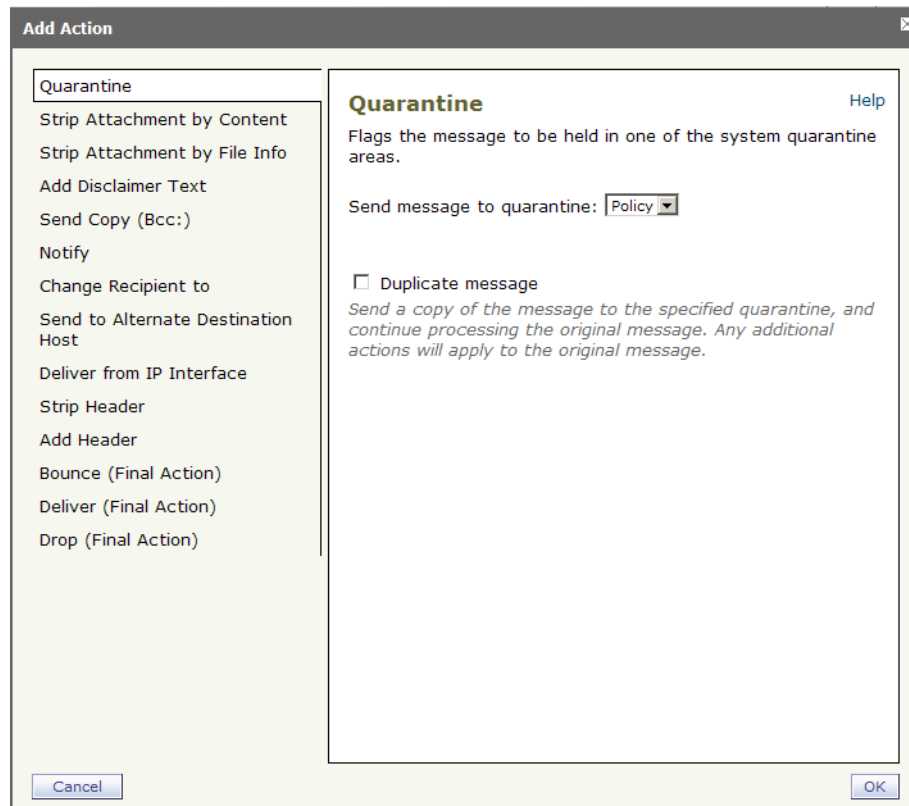


Table 6-2 Content Filter Actions

Action	Description
Quarantine	<p>Quarantine. Flags the message to be held in one of the system quarantine areas.</p> <p>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.</p>
Strip Attachment by Content	<p>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.</p> <p>Contains smart identifier. Drops all attachments on a message that contains the specified smart identifier.</p> <p>Attachment contains terms in the content dictionary. Does the attachment contain any of the regular expressions or terms in the content dictionary named <i><dictionary name></i>?</p> <p>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.</p> <p>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.</p>

Action	Description
Strip Attachment by File Info	<p>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.</p> <p>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.</p> <p>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. For more information, see “Attachment Filtering” in the <i>AsyncOS for Email Advanced User Guide</i>.</p> <p>MIME type. Drops all attachments on messages that have a given MIME type.</p> <p>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.</p>
Add Disclaimer Text	<p>Above. Add disclaimer above message (heading).</p> <p>Below. Add disclaimer below message (footer).</p> <p>Note: You must have <i>already</i> created disclaimer text in order to use this content filter action. For instructions on creating disclaimer text, see Disclaimer Text, page 4-57.</p>
Send Copy (Bcc:)	<p>Email addresses. Copies the message anonymously to the specified recipients.</p> <p>Subject. Add a subject for the copied message.</p> <p>Return path (optional). Specify a return path.</p> <p>Alternate mail host (optional). Specify an alternate mail host.</p>

Action	Description
Notify	<p>Notify. Reports this message to the specified recipients. You can optionally notify the sender and recipients.</p> <p>Subject. Add a subject for the copied message.</p> <p>Return path (optional). Specify a return path.</p> <p>Use template. Select a template from the templates you created.</p> <p>You create a notification template by going to Mail Policies > Text resources. Click Add Text Resource and enter a name and type for your template. Then enter your text and click Submit.</p> <p>Include original message as an attachment. Adds the original message as an attachment.</p>
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.
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.
Add Header	<p>Header name. Inserts a header into the message before delivering.</p> <p>Header value. Inserts a value for the header into the message before delivering.</p>
Bounce (Final Action)	Sends the message back to the sender.
Deliver (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) or quarantine.
Drop (Final Action)	Drops and discards the message.

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 Blocker appliance supports the following set of action variables

Table 6-3 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.
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.

Variable	Syntax	Description
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 Blocker 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 Blocker appliance.
Remote Host Address	\$remotehost	Replaced by the hostname of the system that sent the message to the Blocker 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.
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.

Understanding Threshold Scoring in Content Filtering

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 Blocker 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.

Threshold Scoring for Message Bodies and Attachments

An email message may be composed of multiple parts. When you specify threshold values for filter rules that search for patterns in the message body or attachments, the Blocker counts the number of matches in the message parts and attachments to determine the threshold “score.” Unless the message filter specifies a specific MIME part (such as the attachment-contains filter rule), the Blocker will total the matches found in all parts of the message to determine if the matches total the threshold value. For example, you have a body-contains filter with a threshold of 2. You receive a message in which the body contains one match, and the attachment contains one match. When the Blocker scores this message, it totals the two matches and determines that the threshold score has been met.

Similarly, if you have multiple attachments, Blocker totals the scores for each attachment to determine the score for matches. For example, you have an attachment-contains filter rule with a threshold of 3. You receive a message with two attachments, and each attachment contains two matches. The Blocker would score this message with four matches and determine that the threshold score has been met.

Threshold Scoring Multipart/Alternative MIME Parts

To avoid duplicate counting, if there are two representatives of the same content (plain text and HTML), the Blocker does not total the matches from the duplicate parts. Instead, it compares the matches in each part and selects the highest value. The Blocker would then add this value to the scores from other parts of the multipart message to create a total score.

How Do I Create a List of Terms to Search For?

When you use content filters, you may want to create a list of terms to search for. For example, if your company policy prohibits use of profanity, you could create a list of profane words to search for and then create a content filter that uses this list to search for profanity on outgoing messages. This list of terms is called a *content dictionary*.

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 your list of profane words, and then use a filter rule to scan messages that contain words in the list. When the terms are found, you can configure the filter to drop, archive, or quarantine the message. You can define a total of 32 content dictionaries using the GUI (**Mail Policies > Dictionaries**). You can create, delete, and view dictionaries; add and delete entries from a dictionary; and import and export entire dictionaries.

For each term, you specify a “weight,” so that certain terms can trigger filter conditions more easily. When the Blocker 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. The Blocker 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 be useful for policy enforcement.

The content dictionary feature also includes, by default, the following text files located in the configuration directory of the appliance:

- GLBA-Dictionary.txt
- HIPAA-Dictionary.txt
- PCI-Dictionary.txt

- SOX-Dictionary.txt
- 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.

Adding Dictionaries

To create a new dictionary:

1. Click **Mail Policies > Dictionaries**.
2. Click **Add Dictionary** on the Dictionaries page. The Add Dictionary page is displayed:
3. Type a name for the dictionary.
4. Specify whether to match whole words only by marking the checkbox next to **Match Whole Words Only**.
5. Specify whether to perform case-sensitive searches. See [Matching Case-Sensitive Words](#) for more information.
6. 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.
7. Enter new dictionary entries into the list of terms.
8. 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.
9. Click **Add**.
10. 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. IronPort 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 the Blocker 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.

Dictionary Entry Syntax

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.

Deleting Dictionaries

To delete a dictionary:

1. Click the trash can icon next to the dictionary to delete in the dictionary listing. A confirmation message is displayed.
2. The confirmation message lists any filters that are currently referencing the dictionary.
3. Click **Delete** to delete the dictionary.
4. Submit and commit your changes.

Any content filters that reference the deleted dictionary are left enabled, but will evaluate to false.

Importing and Exporting Dictionaries

To import or import a dictionary:

1. Click **Mail Policies > Dictionaries**.
2. Click **Import** or **Export Dictionary**.
3. You can import from your local computer or configuration file.
OR
You can export to your local computer or configuration file.
4. Select a file to import or export.
5. Select the location to import or export from.



Note

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

6. Select the default weight to use for dictionary terms. The Blocker will assign a default weight to any terms with unspecified weights. You can edit the weights after importing the file.
7. Select an encoding.
8. If you are importing the file, you can edit it or rename it before you import it.

9. Submit and commit your changes.



CHAPTER 7

Monitoring Your Blocker

Revised: October 8, 2009 OL-20859-01

Monitoring Your Blocker: Overview

When your Blocker is up and running, there will be times when you will want to know in more detail how your appliance is performing. You may need to find out what happened to an email sent to one of your mail users. The best tool for this is the **Message Tracking** application which can search for and track messages using a variety of different parameters, such as the sender, subject, etc. This tool is very useful when you are trying to find a missing email. For more information, see [How Can I Find Out What Happened to An Email Sent to My Blocker?](#), page 7-94.

Or, you may just want to get a general picture of the volume of traffic, the number of spam or virus emails blocked, or the volume of messages stopped by the Senderbase Reputation Service (SBRs). To get a general overview of the activity on your Blocker, the best tool is the **Email Security Monitor**. You can go to the overview page, which shows a number of graphs depicting traffic and system performance, or you can run a number of different reports to answer more specific questions. For information about the Email Security Monitor, see [How can I View the Activity on My Blocker?](#), page 7-95

Finally, you can configure logs which will log critical information about your Blocker appliance. A log subscription associates a log type with a name, logging level, and other constraints such as size and destination information; multiple subscriptions for the same log type are permitted. Instructions for enabling logging are listed in [Logging](#), page 7-103.

Chapter Concepts/Glossary

The following concepts are discussed or introduced in this chapter. Some of these terms are specific to the Blocker appliance, while others are networking terms you will need to know in order to configure your Blocker appliance. It's a good idea to review these terms before you begin.

- **Message Tracking.** You can use Message Tracking to find out more information about a specific message as it passes through your Blocker. Message tracking follows the path of an email from when it enters your Blocker until it is delivered to an email user, sent to a quarantine, or blocked.
- **Email Security Monitor.** The email Security Monitor contains a variety of reports that can give you a detailed picture of various aspects of your Blocker, such as trends in your incoming mail, trends in virus and spam email, system health, reports on particular users, and mail delivery status. You can use these reports to get a general sense of the activity on your Blocker or to track down a specific problem.

- **Email Pipeline.** As an email is processed through the Blocker, different actions are taken on the mail in a particular order. The order of the actions is considered the email “pipeline.” It’s important to be aware of this order because an action that occurs early in the pipeline may prevent an action later in the pipeline from taking place. For example, one of the first actions the Blocker performs on a message is to see if an incoming mail matches a HAT entry (The HAT maintains a set of rules that control incoming connections from email senders for a listener). If the given message meets a HAT criteria for being rejected, then the message may not even enter the work queue.

Before you Begin

In order to monitor your email, it’s a good idea to understand the different types of actions that may occur to messages as they pass through your Blocker appliance. This is especially important because there are a number of actions you may configure which can cause a message to be rejected, bounced, or sent to a quarantine. A good way to familiarize yourself with the email pipeline is to review “Understanding the Email Pipeline” on page 40. This section gives you an overview of common actions that are performed on incoming email and where these actions are configured.

How Can I Find Out What Happened to An Email Sent to My Blocker?

The message tracking service makes it easy to find the status of messages processed by the Blocker, and you can quickly resolve help desk calls by determining the exact location of a message. For example, suppose a mail user expected an important email regarding a purchase order from a colleague at another company. Finally, she calls the colleague only to find out that he sent her the critical email hours earlier. You can track down this missing email by the envelope sender (the colleague at another company), the subject header information (such as the purchase order number), the time frame (the last day, for example), sender IP address (by looking up the address of the sending company).

You can also track emails by event information or message IDs. Event information includes mails that match specified events, such as messages flagged as virus positive, spam positive, or suspected spam, and messages that were delivered, hard bounced or soft bounced. *Unlike most conditions that you add to a tracking query, events are added with an “OR” operator.* When you search by message ID, you find emails by identifying the SMTP “Message-ID:” header or the Blocker message ID (MID).

You can use message tracking to determine if a particular message was delivered, found to contain a virus, or placed in a spam quarantine — or if it is located somewhere else in the mail stream.



Note

Another way to troubleshoot mail flow is to use the trace page. You can use **System Administration > Trace** page to debug the flow of messages through the system by emulating sending a test message. The Trace page emulates a message as being accepted by a listener and prints a summary of features that would have been “triggered” or affected by the current configuration (including uncommitted changes) of the system.

Running a Search Query

To run a search query:

1. On the **Monitor > Message Tracking** page, complete the desired search fields.

You do not need to complete every field. Except for the Message Event options, the query is an “AND” search.

2. Click **Search** to submit the query. The query results are displayed at the bottom of the page. Each row corresponds to an email message.
3. Review your query results. To view detailed information about a particular email message, including the message header information and processing details, click the **Show Details** link. A new browser window opens with the message details.

**Note**

Tracking does not support wildcard characters or regular expressions. Tracking searches are not case sensitive.

Narrowing the Result Set

After you run a query, you can narrow the result set by clicking a value within a row. Clicking a value adds the parameter value as a condition in the search. For example, if the query results include messages from multiple dates, you can click a particular date within a row to show only messages that were received on that date.

Disabling Message Tracking

Message tracking is enabled by default. To disable message tracking:

1. On **Services > Message Tracking**, click the **Edit Settings** button.
2. Clear the **Enable Message Tracking Service** check box.
3. Submit and commit your changes.

How can I View the Activity on My Blocker?

You use Email Security Monitor to monitor domains that send mail to your company as well as checking on your system health, reviewing trends in viruses and spam sent to your company, and searching for potential company policy violations. You can monitor, sort, analyze, and classify the “mail flow” of your appliance and differentiate between high-volume senders of legitimate mail and potential “spammers” (senders of high-volume, unsolicited commercial email) or virus senders. These pages can also help you troubleshoot inbound connections to the system.

You can generate a printer-friendly formatted .pdf version of any of the Email Security Monitor pages by clicking on the Printable PDF link at the top-right of the page. You can export graphs and other data to CSV (comma separated values) format via the Export link.

What Kind of Items Can I Search for in the Email Security Monitor?

Many of the Email Security Monitor pages include a search form.

You can search for the following types of items:

- **IP Address:** The numerical label that is assigned to devices participating in a computer network.

- **domain:** Typically, a domain name is used to provide easily recognizable names to numerically addressed Internet resources (IP addresses).
- **network owner:** The registrant of the domain name.
- **internal users:** Your internal mail users.
- **destination domain:** the domain that an email is being sent to.
- **internal sender domain:** The domain associated with an internal sender.
- **internal sender IP address:** The IP address associated with mail sent by your internal sender.
- **outgoing domain deliver status:** The status of emails sent to a particular domain.

For domain, network owner, and internal user searches, choose whether to exactly match the search text or look for items starting with the entered text (for instance, starts with “ex” will match “example.com”).

For IP address searches, the entered text is always interpreted as the beginning of up to four IP octets in dotted decimal format. For instance, “17” will search in the range 17.0.0.0 through 17.255.255.255, so it will match 17.0.0.1 but not 172.0.0.1. For an exact match search, simply enter all four octets. IP address searches also support CIDR format (17.16.0.0/12).

All searches are bounded by the time range currently selected on the page.

Where Can I View a Snapshot of the Activity on my Blocker?

The Email Security Monitor page provides a snapshot summary of the message activity on your Blocker—including incoming an email overview, a quarantine summary, and graphs to give you a visual representation of the state of your email. You can use this page to quickly find out information about the email that flows through your mail server.

On the Overview page, you can:

- View a mail trend graph of all mail “flowing” into or out of your gateway.
- View a graph showing the number of attempted messages, messages stopped by reputation filtering (SBRs), messages with invalid recipients, messages marked as spam, messages marked as virus positive, and clean messages.
- View the summary of the system status and local quarantines.

The Overview page is divided into two sections: System Overview and Incoming and Outgoing Mail graphs and summary.

System Overview

The System Overview section of the Overview page serves as a system dashboard, providing details about the appliance including system and work queue status, quarantine status, and virus outbreak activity.

Status

This section provides an overview of the current state of the appliance and inbound mail processing.

System Status: One of the following states:

- **Online.** The Blocker is receiving mail, processing mail in the work queue, and delivering mail.

- **Resource Conservation.** Conservation occurs when the system is low on available RAM or mail queue space. The Blocker will try to prevent the situation from getting worse by inserting delays into the SMTP conversation which slows down the rate at which new mail is accepted. In an extreme case, the ESA may refuse to accept new mail.
- **Delivery Suspended.** The Blocker will accept new mail and process it through the work queue. The mail will be queued up, but it will not be delivered.
- **Receiving Suspended.** The listeners are shut down, so the Blocker won't respond to any incoming SMTP connections at all.
- **Work Queue Paused.** The Blocker accepts new mail but it will sit in the work queue; mail that has exited the work queue will be delivered.
- **Offline.** The listeners are shut down, so the Blocker won't respond to any incoming SMTP connections at all. If there is mail in the work queue, it will be queued for delivery, but delivery will not take place.

Incoming Messages: The average rate of incoming mail per hour. This projection is based on the number of incoming messages for the last 15 minutes (so if you have received 4 messages in the last 15 minutes, the hourly rate would be 16).

Work Queue: The number of messages awaiting processing in the work queue.

Click the System Status Details link to navigate to the System Status page.

System Quarantines

This section displays information about the top three quarantines by disk usage on the appliance, including the name of the quarantine, how full the quarantine is (disk space), and the number of messages currently in the quarantine.

Click the Local Quarantines link to navigate to the Local Quarantines page.

Virus Threat Level

This section shows the Virus Outbreak status as reported by the Threat Operations Center (TOC). For example, Figure 6-1 shows that a virus outbreak has been identified in the last 24 hours. Also shown is the status of the Outbreak quarantine, including how full it is (disk space) and the number of messages in the quarantine. The Outbreak quarantine is only displayed if you have enabled the Virus Outbreak Filters feature on your appliance.

Incoming and Outgoing Summary and Graph

The Incoming and Outgoing summary sections provide access to real-time activity of all mail activity on your system and is comprised of the Incoming and Outgoing Mail Graphs and Mail Summaries. You can select the time frame on which to report via the Time Range menu. You can select Hour, Day, Week, or Month. The time range you select is used throughout all of the Email Security Monitor pages.

How is Email Categorized?

Messages reported in the Overview and Incoming Mail pages are categorized as follows:

- **Stopped by Reputation Filtering:** All connections blocked by HAT policies multiplied by a fixed multiplier, plus all recipients blocked by recipient throttling.
- **Invalid Recipients:** All recipients rejected by conversational LDAP rejection plus all RAT rejections.

- **Spam Messages Detected:** The total count of messages detected by the anti-spam scanning engine as positive or suspect and also those that were both spam and virus positive.



Note If your spam engine detects a message that is both spam and virus positive, it is counted as a spam email. However, the Virus Types page will also show that the message is virus-positive.

- **Virus Messages Detected:** The total count and percentage of messages detected as virus positive and not also spam.



Note If you have configured your anti-virus settings to deliver unscannable or encrypted messages, these messages will be counted as clean messages and not virus positive. Otherwise, the messages are counted as virus positive.

- **Stopped by Content Filter:** The total count of messages that were stopped by a content filter.
- **Clean Messages Accepted:** Mail that is accepted and is deemed to be virus and spam free — the most accurate representation of clean messages accepted when taking per-recipient scanning actions (such as splintered messages being processed by separate mail policies) into account. However, because messages that are marked as spam or virus positive and still delivered are not counted, the actual number of messages delivered may differ from the clean message count.



Note Messages that match a message filter and are not dropped or bounced by the filter are treated as clean. Messages dropped or bounced by a message filter are not counted in the totals.

Reporting Pages Populated with Data: Sender Profile Pages

If you clicked a sender in the Incoming Mail Summary table listed in the Incoming Mail page, the resulting Sender Profile page is listed with data for the particular IP address, domain, or organization (network owner). Sender Profile pages show detailed information for the sender. You can access a Sender Profile page for any network owner, domain, or IP address by clicking on the specified item in the Incoming Mail or other Sender Profile pages. Network owners are entities that contain domains; domains are entities that contain IP addresses.

Where Can I Find Details about Incoming Mail?

The Incoming Mail page provides access to real-time activity of all public listeners configured on your system and is comprised of two main sections: the mail trend graphs summarizing the top domains received (by total threat messages and by total clean messages) and the Incoming Mail Details listing.

Notes on Time Ranges in the Mail Trend Graph

The Email Security Monitor feature constantly records data about the mail flowing into your gateway. The data are updated every 60 seconds, but the display shown is delayed by 120 seconds behind the current system time. You can “zoom out” from the day view to the week, day, and month views of the same data. Because the data is monitored in real time, information is periodically updated and summarized in the database.

Incoming Mail Details Listing

The top senders which have connected to public listeners of the appliance are listed in the External Domains Received listing table at the bottom of the Incoming Mail page, based on the view selected. Click the column headings to sort the data.

The system acquires and verifies the validity of the remote host's IP address (that is, the domain) by performing a double DNS lookup.

The Sender Detail listing has two views, Summary and All.

The *Summary* view shows the total number of attempted messages for each sender, and includes a breakdown by category (the same categories as the Incoming Mail Summary graph on the Overview page: number of clean messages, stopped by reputation filtering, invalid recipients, spam detected, virus detected, and stopped by content filter).

The *All* view shows the connection information (Accepted, Rejected) for senders as well as the breakdown by category. An additional column, Stopped by Recipient Throttling, lists throttled recipients.

Sort the listing by clicking the column header links. The sorting order is retained when you switch between the summary and all views, regardless of whether or not the sorted column exists in both views. In other words, if you sort the summary listing by "Total Attempted" and then switch to the All view, the data will retain its sorting.

Where Can I Find Information about Where My Company Sends Mail?

The Outgoing Destinations page provides information about the domains your company sends mail to. The page consists of two sections. The top half of the page consists of graphs depicting the top destinations by outgoing threat messages and top destinations by outgoing clean messages on the top half of the page. The bottom half of the page displays a chart showing all the columns sorted by total recipients (default setting).

You can select a time range on which to report (hour, day, week, or month). As with all reports, you can export the data for the graphs or the details listing to CSV format via the Export link.

The Outgoing Destinations page can be used to answer the following types of questions:

- What domains is the Blocker appliance sending mail to?
- How much mail is sent to each domain?
- How much of that mail is clean, spam-positive, virus-positive, or stopped by a content filter?
- How many messages are delivered and how many messages are hard-bounced by the destination server?

Where Can I find Details about the Mail Sent from My Mail Users?

The Outgoing Senders page provides information about the quantity and type of mail being sent from IP addresses and domains in your network. You can view the results by domain or IP address when you view this page. You might want to view the results by domain if you want to see what volume of mail is being sent by each domain, or you might want to view the results by IP address if you want to see which IP addresses are sending the most virus messages or triggering content filters.

The page consists of two sections. On the left side of the page is a graph depicting the top senders by total threat messages. Total threat messages include messages that are spam or virus positive or triggered a content filter. On the right side of the page is a graph displaying top senders by clean messages on the top half of the page. The bottom half of the page displays a chart showing all the columns sorted by total messages (default setting).

This page does not display information about message delivery. Delivery information, such as how many messages from a particular domain were bounced can be tracked using the Delivery Status page.

You can select a time range on which to report (hour, day, week, or month). As with all reports, you can export the data for the graphs or the details listing to CSV format via the Export link.

The Outgoing Senders page can be used to answer the following types of questions:

- Which IP addresses are sending the most virus or spam positive email?
- Which IP addresses trigger content filters the most frequently?
- Which domains are sending the most mail?

Where Can I Find Out if My Mail is Being Delivered Correctly?

If you suspect delivery problems to a specific recipient domain or if you want to gather information on a Virtual Gateway address, the **Monitor > Delivery Status** Page provides monitoring information about email operations relating to a specific recipient domain.

This page displays a list of the top 20, 50, or 100 recipient domains for messages delivered by the system within the last three hours. You can sort by latest host status, active recipients (the default), connections out, delivered recipients, soft bounced events, and hard bounced recipients by clicking the links in the column heading for each statistic.

To search for a specific domain, type the name of the domain in the Domain Name: field and click Search.

To drill down on a domain shown, click the domain name link.

The results are shown in an Delivery Status Details Page.

Any activity for a recipient domain results in that domain being “active” and thus present in the overview page. For example, if mail remains in the outbound queue due to delivery problems, that recipient domain continues to be listed in the outgoing mail overview.

Where Can I Drill Down on Information about a Specific Mail User?

The Internal User detail page shows detailed information about the specified user, including a breakdown of incoming and outgoing messages showing the number of messages in each category (spam detected, virus detected, stopped by content filter, and clean). Incoming and outgoing content filter matches are also shown.

Click on a content filter name to view detailed information for that filter in the corresponding content filter information page. You can use this method to get a list of users who also sent or received mail that matched that particular content filter.

How Can I Find Out Which Content Filters Have Triggered the Most Matches?

The Content Filters page shows information about the top content filter matches (which content filter had the most matching messages) in two forms: a bar chart and a listing. Using the Content Filters page, you can review your corporate policies on a per-content filter or per-user basis and answer questions like:

Which content filter is being triggered the most by incoming or outgoing mail?

Who are the top users sending or receiving mail that is triggering a particular content filter?

You can click the name of the content filter in the listing to view more information about that filter on the Content Filter detail page.

Where Can I Find More Information on the Viruses Detected By My Blocker?

The Virus Types page provides an overview of the viruses entering and being sent from your network. The Virus Types page displays the viruses that have been detected by the virus scanning engines running on your Blocker. You might want to use this report to take a specific action against a particular virus. For example, if you see that you are receiving a high volume of a viruses known to be embedded in PDF files, you might want to create a filter action to quarantine messages with PDF attachments.

Where Can I Track Information about My TLS Connections?

The TLS Connections pages shows the overall usage of TLS connections for sent and received mail. The report also shows details for each domain sending mail using TLS connections.

The TLS Connections page can be used to determine the following information:

- Overall, what portion of incoming and outgoing connections use TLS?
- What partners do I have successful TLS connections with?
- What partners do I have unsuccessful TLS connections with?
- What partners have issue with their TLS certificates?
- What percent of overall mail with a partner uses TLS?

Where Can I Find Out How Well the Blocker Can Handle My Mail Volume?

The System Capacity page provides a detailed representation of the system load, including messages in the work queue, average time spent in the work queue, incoming and outgoing messages (volume, size, and number), overall CPU usage, CPU usage by function, and memory page swapping information.

- The system capacity page can be used to determine the following information:
- Identify when Blocker appliance is exceeding recommended capacity and configuration optimization or additional appliances are needed.
- Identify historical trends in system behavior which point to upcoming capacity issues.
- Identify which part of the system is using the most resources to assist with troubleshooting.

Where Can I Find Out More About My System's Health?

The System Status page provides a detailed representation of all real-time mail and DNS activity for the system. The System Status page is comprised of four sections: System Status, Gauges, Rates, and Counters.

System Status

The system status section shows Mail System Status and Version Information.

Mail System Status

The Mail System Status section includes:

- System Status
- The last time the status was reported.
- The uptime for the appliance.
- The oldest message in the system, including messages that have not yet been queued for delivery.

Version Information

The Version Information section includes:

- The Blocker appliance model name.
- The version and build date of the operating system installed.
- The installation date of the operating system.
- The serial number of the system to which you are connected.

Gauges

The Gauges section shows queue and resource utilization.

- Mail Processing Queue
- Active Recipients in Queue
- Queue Space
- CPU Utilization

Rates

The Rates section shows rate handling for recipients.

- Mail Handling Rates
- Completion Rates

Counters

You can reset the cumulative email monitoring counters for system statistics and view the last time the counters were reset. The reset affects system counters as well as per-domain counters. The reset does not affect the counters on messages in the delivery queue related to retry schedules.

Logging

Logs are a compact, efficient method of gathering critical information about the email operations on your Blocker. These logs record information regarding activity on your Blocker appliance. The information will vary depending upon the log you view, for example, Bounce logs or Delivery logs.

Most logs are recorded in plain text (ASCII) format; however, delivery logs are formatted in binary for resource efficiency. The ASCII text information is readable in any text editor.

Log Rollover and Transfer Schedule

Log subscriptions create and transfer (rollover) log files based on the first user-specified limit reached: maximum file size or maximum time. Log subscriptions based on the FTP poll transfer mechanism will create files and store them in the FTP directory on the Blocker appliance until they are retrieved or until the system needs more space for log files.

Log Types

A log subscription associates a log type with a name, logging level, and other constraints such as size and destination information; multiple subscriptions for the same log type are permitted. The log type indicates what information will be recorded within the generated log such as message data, system statistics, binary or textual data. You select the log type when creating a log subscription.

The Blocker generates the following log types:

The following log types are available by default:

Table 7-1 Log Types

Log Name	Description
antispan	Anti-Span Logs -Anti-spam logs record the status of the anti-spam scanning feature of your system, including the status on receiving updates of the latest anti-spam rules. Also, any logs related to the Context Adaptive Scanning Engine are logged here.
antivirus	Anti-Virus Logs -AntiVirus logs record the status of the anti-virus scanning feature of your system, including the status on receiving updates of the latest anti-virus identity files.
asarchive	Anti-Span Archive - If you enabled an Anti-Span scanning feature, messages that are scanned and associated with the “archive message” action are archived here. The format is an mbox-format log file.
authentication	Authentication Logs - The authentication log records successful user logins and unsuccessful login attempts.
avarchive	Anti-Virus Archive - If you enabled an anti-virus engine, messages that are scanned and associated with the “archive message” action are archived here. The format is an mbox-format log file.

Table 7-1 Log Types (continued)

Log Name	Description
bounce	Bounce Logs -Bounce logs record information about bounced recipients. The information recorded for each bounced recipient includes: the message ID, the recipient ID, the Envelope From address, the Envelope To address, the reason for the recipient bounce, and the response code from the recipient host. In addition, you can choose to log a fixed amount of each bounced recipient message. This amount is defined in bytes and the default is zero.
cli-logs	CLI Audit Logs -The CLI audit logs record all CLI activity on the system.
error_logs	Error logs- mail logs record information regarding the errors on the email system.
euq-logs	IronPort Spam Quarantine Logs - Spam Quarantine logs record actions associated with the IronPort Spam Quarantine processes.
euqgui_logs	IronPort Spam Quarantine GUI Logs - Spam Quarantine logs record actions associated with the Spam Quarantine including configuration via the GUI, end user authentication, and end user actions (releasing email, etc.).
ftpd_logs	FTP Server Logs -FTP logs record information about the FTP services enabled on the interface. Connection details and user activity are recorded.
gui_logs	HTTP Logs - HTTP logs record information about the HTTP and/or secure HTTP services enabled on the interface. Because the graphical user interface (GUI) is accessed via HTTP, the HTTP logs are ostensibly the GUI equivalent of the CLI Audit logs. Session data (new session, session expired) and pages accessed in the GUI are recorded.
lmail_logs	Text Mail Logs -Text mail logs record information regarding the operations of the email system. For example, message receiving, message delivery attempts, open and closed connections, bounces and others.
reportd_logs	Reporting Logs -Reporting logs record actions associated with the processes of the centralized reporting service.
reportqueryd_logs	Reporting Query Logs -Reporting query logs record actions associated with the reporting queries that are run on the appliance.
scanning	Scanning Logs -The scanning log contains all LOG and COMMON messages. This is typically application faults, alert sent, alert failed, and log error messages.
slbl_logs	Safelist/Blocklist logs- Safelist/blocklist logs record data about the safelist/blocklist settings and database.
snmp_logs	SMTP Conversation Logs - The SMTP conversation log records all parts of incoming and outgoing SMTP conversations.

Table 7-1 Log Types (continued)

Log Name	Description
sntpd_logs	NTP Logs - NTP logs record the conversation between the appliance and any NTP (Network Time Protocol) servers configured.
status	Status Logs - This log file records system statistics found in the CLI status commands, including <code>status detail</code> and <code>dnsstatus</code> . The period of recording is set using the <code>setup</code> subcommand in <code>logconfig</code> . Each counter or rate reported in status logs is the value since the last time the counter was reset.
system_logs	System Logs - System logs record the following: boot information, DNS status information, and committed comments. System logs are useful for troubleshooting the basic state of the appliance.
trackerd_logs	Tracking Logs - Tracking logs record actions associated with the processes of the tracking service. Tracking logs are a subset of the mail logs.
updater_logs	Updater Logs - The updater log records events related to updates for system services, such as McAfee Anti-Virus definition updates.

Retrieving Your Mail Logs

Log files can be retrieved based upon one of the following file transfer protocols. You set the protocol while creating or editing the log subscription in the GUI during the log subscription process.

Table 7-2 Log Transfer Protocols

FTP Poll	This method involves a remote FTP client accessing the Blocker to retrieve log files using an admin or operator user's username and password. When configuring a log subscription to use the FTP poll method, you must supply the maximum number of log files to keep on hand. When the maximum number is reached, the system deletes the oldest file.
FTP Push	This method periodically pushes log files to an FTP server on a remote computer. The subscription requires a username, password, and destination directory on the remote computer. Log files are transferred based on a rollover schedule set by you.

Table 7-2 *Log Transfer Protocols (continued)*

SCP Push	This method periodically pushes log files to an SCP server on a remote computer. This method requires an SSH SCP server on a remote computer using the SSH1 or SSH2 protocol. The subscription requires a username, SSH key, and destination directory on the remote computer. Log files are transferred based on a rollover schedule set by you.
Syslog Push	This method sends log messages to a remote syslog server. This method conforms to RFC 3164. You must submit a hostname for the syslog server and choose to use either UDP or TCP for log transmission. The port used is 514. A facility can be selected for the log; however, a default for the log type is pre-selected in the dropdown menu. Only text-based logs can be transferred using syslog push.

Troubleshooting Tips

This section addresses common problems that occur when installing the Blocker:

- **Problem:** On the Email Security Monitor overview page, I don't see any messages marked as spam. But when I look at the Virus Types page, I show that some viruses were found. Why do these reports differ?

Solution: When a message that is both spam and virus positive is caught by the spam engine, the engine flags the message as spam (by default the spam engine checks for spam first, and marks a message as spam even if it is spam and virus positive). Therefore when viewing reports on incoming mail and spam filter results, a virus-positive message would not appear. However, the Virus Types page is designed to detect virus messages regardless of how they were categorized by the spam filter. This means that the reports may sometimes show some discrepancies.



CHAPTER 8

Advanced Configuration

Revised: October 8, 2009 OL-20859-01

Advanced Configuration: Overview

The Blocker is a robust appliance with a variety of advanced options that are not discussed in this guide. This guide is intended to cover most common scenarios. For these advanced topics, see the PDF versions of the *AsyncOS for Email Configuration Guide*, the *AsyncOS for Email Daily Configuration Guide*, and the *AsyncOS for Email Advanced Configuration Guide* included in your documentation CD.

Listing of Advanced Topics

The following section includes a brief description of some advanced functionality you can configure.

- **LDAP Configuration.** This guide includes the steps for creating an LDAP profile. However, there is more functionality available than simply using your LDAP server for authentication. You can work with your LDAP profile to perform the following actions:
 - **routing-** route messages to the appropriate address and/or mail host based upon the information available in LDAP directories on your network.
 - **Masquerading-** You can masquerade Envelope Senders (for outgoing mail) and message headers (for incoming mail, such as To:, Reply To:, From: or CC:).
 - **Group Queries** - perform actions on messages based on the groups in the LDAP directory.
 - **Domain-based queries-** domain-based queries allow the Blocker to perform different queries for different domains on a single listener.
 - **Chain queries** - A chain query enables the Blocker to perform a series of queries in sequence.
 - **Directory harvest prevention-** The Blocker can combat directory harvest attacks using your LDAP directories. You can configure directory harvest prevention during the SMTP conversation or within the work queue.
 - **SMTP Authentication** - The Blocker provides support for SMTP authentication.
 - **External authentication** - You can configure your Blocker to use your LDAP directory to authenticate users logging in to the Blocker.
 - **Spam quarantine end-user authentication** - You can configure your Blocker to validate users when they log in to the end-user quarantine.

- **Spam quarantine alias consolidation** - If you use email notifications for spam, this query consolidates the end-user aliases so that end-users do not receive quarantine notices for each aliased email address.

For more information, see the *AsyncOS for Email Advanced Configuration Guide*.

- **SMTP Routes.** SMTP Routes allow you to redirect all email for a particular domain to a different mail exchange (MX) host. For example, you could make a mapping from example.com to groupware.example.com. This mapping causes any email with @example.com in the Envelope Recipient address to go instead to groupware.example.com. The system performs an “MX” lookup on groupware.example.com, and then performs an “A” lookup on the host, just like a normal email delivery. This alternate MX host does not need to be listed in DNS MX records and it does not even need to be a member of the domain whose email is being redirected. The Blocker allows up to ten thousand (10,000) SMTP Route mappings to be configured for your Blocker appliance. For more information, see the *AsyncOS for Email Advanced Configuration Guide*.
- **Domain Keys and DKIM Signing.** Blocker supports DomainKeys and DKIM authentication to prevent email forgery. DomainKeys and DKIM are mechanisms used to verify that the source of the email and the contents of the message were not altered during transit. DKIM is an enhanced protocol that combines DomainKeys specification with aspects of Identified Internet Mail to create an enhanced protocol called DomainKeys Identified Mail (DKIM). DomainKeys and DKIM consist of two main parts: signing and verification. The current version of Blocker supports the “signing” half of the process for DomainKeys, and it supports both signing and verification for DKIM. You can also enable bounce and delay messages to use DomainKeys and DKIM signing. For more information, see the *AsyncOS for Email Advanced Configuration Guide*.
- **Logging.** An important feature within the Blocker is its logging capabilities. The Blocker can generate many types of logs, recording varying types of information. Log files contain the records of regular operations and exceptions from various components of the system. This information can be valuable when monitoring your Blocker appliance as well as when troubleshooting or checking performance. For more information, see the *AsyncOS for Email Advanced Configuration Guide*.
- **Bounce Profiles.** You use the Bounce Profiles page on the Network menu in the GUI to configure how the Blocker handles hard and soft conversational bounces for each listener you create. You create bounce profiles and then apply profiles to each listener via the **Network > Listeners** page (or the listenerconfig command). You can also assign bounce profiles to specific messages using message filters. For more information, see the *AsyncOS for Email Advanced Configuration Guide*.
- **Destination Controls.** For each domain, you can assign a maximum number of connections, outbound messages, and recipients that will never be exceeded by the system in a given time period. This “good neighbor” table is defined through the Destination Controls feature (**Mail Policies > Destination Control**). For more information, see the *AsyncOS for Email Advanced Configuration Guide*.
- **SMTP Authentication.** Blocker provides support for SMTP authentication. SMTP Auth is a mechanism for authenticating clients connected to an SMTP server. For more information, see the *AsyncOS for Email Advanced Configuration Guide*.

The practical use of this mechanism is that users at a given organization are able to send mail using that entity’s mail servers even if they are connecting remotely (e.g. from home or while traveling). Mail User Agents (MUAs) can issue an authentication request (challenge/response) when attempting to send a piece of mail.

Users can also use SMTP authentication for outgoing mail relays. This allows the Blocker appliance to make a secure connection to a relay server in configurations where the appliance is not at the edge of the network.

- **SNMP Monitoring.** The Blocker supports system status monitoring via SNMP (Simple Network Management Protocol). This includes Blocker's Enterprise MIB, ASYNCOS-MAIL-MIB. The ASYNCOS-MAIL-MIB helps administrators better monitor system health. For more information, see the *AsyncOS for Email Advanced Configuration Guide*.
- **Customizing Listeners.** The Blocker supports many options for customizing listeners, such as encrypting SMTP conversations using TLS. For more information, see the *AsyncOS for Email Advanced Configuration Guide*.



GLOSSARY

Revised: October 8, 2009

OL-20859-01

A

- Allowed Hosts** Computers that are allowed to relay email through the Blocker appliance via a private listener. Allowed hosts are defined by their hostnames or IP addresses.
- Anti-Virus** Sophos Anti-Virus scanning engine provide anti-virus protection, detection and disinfection. through virus detection engine 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 marlier 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.
- 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	<p>The act of delivering email messages to recipient domains or internal mail hosts from the Blocker appliance from a specific IP interface. The Blocker 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.</p> <p>You can tailor the configuration of the delivery that the Blocker 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.</p>
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 Blocker appliances. Email Security Manager allows you to manage 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.
False Negative	A spam or virus message that was not detected as such.
False Positive	A message falsely categorized as spam or as containing a virus.

F

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.
---	--

G

Gateway / Enterprise Gateway A gateway SMTP system (usually referred to just as a “gateway”) receives mail from a client system in one transport environment and transmits it to a server system in another transport environment. (Definition from RFC 2821.) In this guide, an Enterprise Gateway is a configuration of an Blocker Blocker that 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.

Greylist A list of suspected senders. The Blocker appliance includes a greylist in the form of the Suspectlist sender group.

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, maximum recipients per hour, the maximum number of recipients per message, the maximum message size, and/or the maximum number of concurrent connections you are willing to accept from a remote host.

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 Blocker appliance — either from the internal systems within your network or from the Internet. Blocker 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.

Blocker 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 Blocker 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 Blocker 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 Blocker 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 Blocker 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.

Q

Queue In the Blocker 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 Brightmail 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 Blocker 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, SBRS 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 Blocker 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.



APPENDIX **A**

User License Agreement

Revised: October 8, 2009 OL-20859-01

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, LL.D., 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 CISCO IRONPORT LLC, 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 Cisco or its reseller (“Agreement”) and the applicable user interface and Cisco’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 Cisco or its reseller, as the case may be.

1.4 “Software” means: (i) Cisco IronPort LLC’s proprietary software licensed by Cisco IronPort LLC to Company along with Cisco IronPort LLC’s hardware products; (ii) any software provided by Cisco IronPort LLC’s third-party licensors that is licensed to Company to be implemented for use with Cisco

IronPort LLC's hardware products; (iii) any other Cisco IronPort LLC software module(s) licensed by Cisco IronPort LLC to Company along with Cisco IronPort LLC'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 Cisco IronPort LLC 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 Cisco IronPort LLC 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 Cisco IronPort LLC 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 Cisco IronPort LLC 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, Cisco IronPort LLC hereby grants to Company a non-exclusive, non-sublicensable, non-transferable, worldwide license during the Term to use the Software only on Cisco IronPort LLC'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 Cisco IronPort LLC, Cisco IronPort LLC'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 Cisco IronPort LLC Privacy Statement at <http://www.IronPort.com/privacy.html>, as the same may be amended from time to time by Cisco IronPort LLC with notice to Company, Company hereby consents and grants to Cisco IronPort LLC 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 Cisco IronPort LLC ("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 Cisco IronPort LLC'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 Cisco IronPort LLC 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 Cisco IronPort LLC or its reseller to Company will remain the exclusive property of Cisco IronPort LLC 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 Cisco IronPort LLC 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 Cisco IronPort LLC 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 Cisco IronPort LLC 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. Cisco IronPort LLC 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 CISCO IRONPORT LLC’S ENTIRE LIABILITY, WILL BE PROMPT CORRECTION OF ANY ERROR OR NONCONFORMITY, PROVIDED THAT THE NONCONFORMITY HAS BEEN REPORTED TO CISCO IRONPORT LLC 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. Cisco IronPort LLC 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 Cisco IronPort LLC 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, CISCO IRONPORT LLC LICENSES THE SOFTWARE AND SERVICES HEREUNDER ON AN “AS IS” BASIS. EXCEPT AS SPECIFICALLY SET FORTH HEREIN, CISCO IRONPORT LLC 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 CISCO IRONPORT LLC 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 Cisco IronPort LLC 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, Cisco IronPort LLC 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 Cisco IronPort LLC or its reseller or destroy all copies of the Software and any other materials or documentation provided to Company by Cisco IronPort LLC 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 Cisco IronPort LLC 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 Cisco IronPort LLC, except that Cisco IronPort LLC may modify the Cisco IronPort LLC 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 Cisco IronPort LLC or a duly authorized representative of Cisco IronPort LLC. 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. CISCO IRONPORT LLC CONTACT INFORMATION. If Company wants to contact Cisco IronPort LLC 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.asw-===}

