



Cisco 6.6 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

Text Part Number: OL-18046-02

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 Spam and Virus Blocker, 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 6.6 Spam and Virus Blocker User Guide
© 2008 Cisco Systems, Inc. All rights reserved.



CONTENTS

Preface vii

How to Use This Guide vii

How This Book is Organized viii

CHAPTER 1

Installing and Administering the Blocker 1-1

Installing the Blocker Overview 1-1

 Before You Begin 1-1

 Register the Blocker Appliance in DNS 1-2

 Network Configuration 1-2

 Configuration Scenarios 1-3

 Segregating Incoming and Outgoing Mail 1-3

 DNS Settings 1-3

 NAT Settings 1-4

 Firewall Settings 1-5

 Running the System Setup Wizard 1-7

 Configuring an Active Directory Server Profile 1-12

 The Blocker Interface 1-13

 Logging In 1-13

 The Commit Changes Button 1-14

 System Administration 1-14

 Shutting Down the Blocker 1-15

 Rebooting the Blocker 1-15

 Open a Support Case 1-15

 Working with Feature Keys 1-16

 The Feature Keys Page 1-16

 Feature Key Settings 1-17

 Editing Update Settings 1-17

 Adding Users 1-18

 Managing the Configuration File 1-19

 Accessing the Appliance 1-20

 Understanding the Email Pipeline 1-20

CHAPTER 2

Accepting Mail 2-23

 Accepting Mail: Overview 2-23

- Configuring the Blocker to Accept Mail (Public Listeners) 2-23
 - IP Interfaces 2-24
 - Incoming Relays 2-24
- Filtering Your Mail Based on the Sender 2-25
 - The Host Access Table (HAT) 2-25
 - Default HAT Entries 2-26
 - Mail Flow Policies: Access Rules and Parameters 2-26
 - HAT Variable Syntax 2-30
 - Using HAT Variables 2-30
- Understanding the Default Sender Groups and Mail Flow Policies 2-30
 - Understanding SBRS (SenderBase Reputation Service) 2-31
 - Implementing SenderBase Reputation Filters 2-32

CHAPTER 3

Mail Routing For Users and Groups 3-35

- Creating Mail Rules for Specific Users , Groups, or Domains 3-35
- Defining the Users You Accept Mail For (Using RAT) 3-36
 - Adding a New RAT Entry 3-37
- Using LDAP to Route Mail for Users and Groups of Users 3-37
 - Configuring the Blocker to work with LDAP 3-38
- Understanding Mail Policies 3-39
 - Incoming vs. Outgoing Messages 3-40
 - Policy Matching 3-40
 - First Match Wins 3-41
 - Examples of Policy Matching 3-41
 - Creating a New Policy 3-42

CHAPTER 4

Sending Mail 4-45

- Sending Mail Overview 4-45
 - Basic Steps for Configuring Outbound Mail 4-45
- Configuring the Listener to Send Mail From Your Network 4-46
 - Creating a Mail Flow Policy for Relaying Mail 4-46
- Adding Disclaimers to Outgoing Mail 4-48
 - Disclaimer Text 4-48
 - Adding Disclaimer Text via a Listener 4-49
- Using Email Tagging to Prevent Bounce Attacks 4-49
 - Configuring Bounce Verification 4-49

CHAPTER 5**Configuring Spam and Virus Protection 5-51**

- Virus and Spam Protection Overview 5-51
- Enabling Virus Protection 5-52
 - Steps to Enable Virus Protection 5-52
 - Editing the Anti-Virus Settings for a Mail Policy 5-52
- Enabling Spam Protection 5-54
 - Steps to Enable Anti-Spam Scanning 5-55
- Configuring the Spam Quarantine 5-55
 - Enabling and Disabling the Spam Quarantine 5-56
 - Disabling the Spam Quarantine 5-56
 - Configuring Settings for the Spam Quarantine 5-56
- Enabling End User Safelists and Blocklists 5-57
- Creating and Maintaining Safelists and Blocklists 5-57
 - Administrator Tasks for Creating and Maintaining Safelists and Blocklists 5-57
- Enabling and Configuring Safelist/Blocklist Settings 5-58
 - Backing Up and Restoring the Safelist/Blocklist Database 5-58
 - Troubleshooting Safelists and Blocklists 5-59
 - End User Tasks for Configuring Safelists and Blocklists 5-59
- Adding Entries to Blocklists 5-61

CHAPTER 6**Enforcing Policies Using Content Filters 6-63**

- Content Filters Overview 6-63
 - To Create a Content Filter 6-63
- Content Filter Conditions 6-64
- Content Filter Actions 6-69
 - Action Variables 6-72

CHAPTER 7**Monitoring Your Email 7-75**

- Email Monitoring Overview 7-75
- Finding Emails Using Message Tracking 7-75
 - Running a Search Query 7-75
 - Narrowing the Result Set 7-76
- Using the Email Security Monitor 7-76
 - Searching and Email Security Monitor 7-76
- Email Security Monitor Overview Page 7-77
 - System Overview 7-77
 - Incoming and Outgoing Summary and Graph 7-78
- Categorizing Email 7-78

- Incoming Mail 7-79
 - Notes on Time Ranges in the Mail Trend Graph 7-79
 - Incoming Mail Details Listing 7-79
- Outgoing Destinations 7-80
- Outgoing Senders 7-80
- The Delivery Status Page 7-80
- Internal User Details 7-81
- The Content Filters Page 7-81
- Virus Types Page 7-81
- TLS Connections Page 7-81
- The System Capacity Page 7-82
- The System Status Page 7-82
- Debugging Mail Flow Using Test Messages: Trace 7-83

CHAPTER 8

Advanced Topics 8-85

- Advanced Topics Overview 8-85
- Listing of Advanced Topics 8-85

GLOSSARY

APPENDIX A

User License Agreement A-7

- License Agreement A-7



Preface

Revised: January 22, 2008, OL-18046-02

About Blocker

This chapter describes the organization of the Blocker Guide. For information about the current release, see the product release notes, which are available on the documentation CD shipped with your Blocker appliance.

You can also view electronic versions of several Blocker 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.

How to Use This Guide

Use this guide as a resource to learn about the features and basic configuration of your Blocker appliance. The topics 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.

This guide is *not* designed to function as an exhaustive reference guide. For each topic discussed in the Blocker User Guide, you may find more detailed information in the following reference guides available in electronic form in the Documentation CD shipped with your Blocker appliance:

- *AsyncOS for Email User Guide*
- *AsyncOS for Email Advanced User Guide*

Where these guides discuss “IronPort appliances” or the “AsyncOS operating system”, you can substitute the terms “Blocker appliances” and “Blocker operating system.” Be aware that not all features listed in these guides are available on the Blocker appliance. For a list of these features, see the “Advanced Topics” chapter.

The guide is distributed electronically as PDF and HTML files. 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:

- **Installing and Administering the Blocker.** This chapter includes the steps needed to install and run your Blocker appliance. It includes information about the System Setup Wizard, network settings, and basic administration tasks, such as shutting down the appliance. See [Installing and Administering the Blocker, page 1](#).
- **Accepting Mail.** This chapter includes information about configuring your Blocker to accept incoming mail. It describes features such as the IP interfaces, Host Access Table (HAT), SenderBase reputation service, that you can configure. See [Accepting Mail, page 23](#).
- **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 35](#).
- **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 51](#).
- **Configuring Spam and Virus Protection.** This chapter includes information about configuring Spam and Virus protection. It describes 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 51](#).
- **Enforcing Policies 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 [Enforcing Policies Using Content Filters, page 63](#).
- **Monitoring Your Email.** This chapter describes the tools available for monitoring your email: Message Tracking and the Email Security Monitor. 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. [Monitoring Your Email, page 75](#).
- **Advanced Topics.** This section describes advanced topics not discussed in this guide. These topics are discussed in the following reference guides: *AsyncOS for Email User Guide*, and the *AsyncOS for Email Advanced User Guide*. See [Advanced Topics, page 85](#).
- **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 Topics, page 85](#).



CHAPTER 1

Installing and Administering the Blocker

Revised: January 22, 2008, OL-18046-02

Installing the Blocker 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 enabled on your system as an Enterprise Gateway (accepting email from the Internet).

Before You Begin

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

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

The Blocker appliance is designed to serve as your SMTP gateway, also known as a mail exchanger or “MX.” In addition to the “hardened” operating system dedicated for Internet messaging, many of the newest features in the Blocker function optimally when the appliance is situated at the first machine with an IP address that is directly accessible to the Internet (that is, it is an external IP address) for sending and receiving email.

The per-recipient reputation filtering, anti-spam, anti-virus are designed to work with a direct flow of messages from the Internet and from your internal network. You can configure the Blocker appliance for policy enforcement for all email traffic to and from your enterprise.

You need to ensure that the Blocker appliance is both accessible via the public Internet and is the “first hop” in your email infrastructure. If you allow another MTA to sit at your network’s perimeter and handle all external connections, then the Blocker appliance will not be able to determine the sender’s IP address. The sender’s IP address is needed to identify and distinguish senders in the Mail Flow Monitor, to query the SenderBase Reputation Service for the sender’s SenderBase Reputation Service Score (SBRs), and to improve the efficacy of the Blocker Anti-Spam feature.

When you use the Blocker appliance as your SMTP gateway:

- The Email Security Monitor feature offers complete visibility into all email traffic for your enterprise from both internal and external senders.
- LDAP queries for routing, aliasing, and masquerading can consolidate your directory infrastructure and provide for simpler updates.

- Familiar tools like alias tables (“Creating Alias Tables” in the *AsyncOS for Email Advanced User Guide*), domain-based routing (“The Domain Map Feature” in the *AsyncOS for Email Advanced User Guide*), and masquerading (“Configuring Masquerading” in the *Blocker AsyncOS for Email Advanced User Guide*) make the transition from Open-Source MTAs easier.

Register the Blocker Appliance in DNS

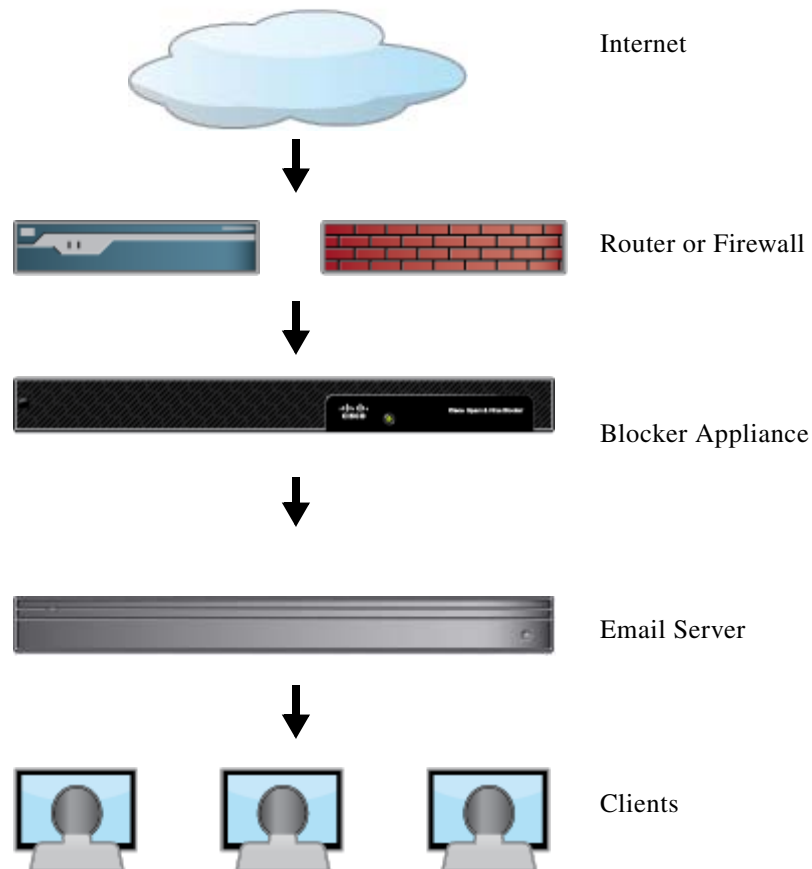
Malicious email senders actively search public DNS records to hunt for new victims. You need to ensure that the Blocker appliance is registered in DNS, if you want to utilize the full capabilities of Blocker Anti-Spam and Sophos Anti-Virus. To register the Blocker appliance in DNS, create an A record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

Network Configuration

You may want to review all features of the appliance prior to installing. The following figure shows the general data flow required when routing mail to your Blocker.

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

Figure 1-1 Network Configuration



In some scenarios, the Blocker appliance resides inside the network “DMZ,” in which case an additional firewall sits between the Blocker appliance and the groupware server. For more information about configuring this scenario, see “Installation” in the *AsyncOS for Email User Guide*.

Configuration Scenarios

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

- **Interfaces** - Only one of the two available Ethernet interfaces on the Blocker appliance is required for most network environments.
- **Public Listener** (incoming email) - The public listener receives connections from many external hosts and directs messages to a limited number of internal groupware servers.
 - Accepts connections from external mail hosts based on settings in the HAT. By default, the HAT is configured to ACCEPT connections from all external mail hosts.
 - Accepts incoming mail only if it is addressed for the local domains specified in the RAT. All other domains are rejected.
 - Relays mail to the appropriate internal groupware server, as defined by SMTP Routes.
- **Private Listener** (outgoing email) - The private listener receives connections from a limited number of internal groupware servers and directs messages to many external mail hosts.
 - Internal groupware servers are configured to route outgoing mail to the Blocker.
 - The Blocker appliance accepts connections from internal groupware servers based on settings in the HAT. By default, the HAT is configured to RELAY connections from all internal mail hosts.

Segregating Incoming and Outgoing Mail

The System Setup Wizard creates an initial configuration of one listener on one logical IP address configured on one physical interface, configured for incoming mail. However, you can add outgoing mail, additional listeners, and segregate incoming from outgoing mail on separate listeners if you wish.

DNS Settings

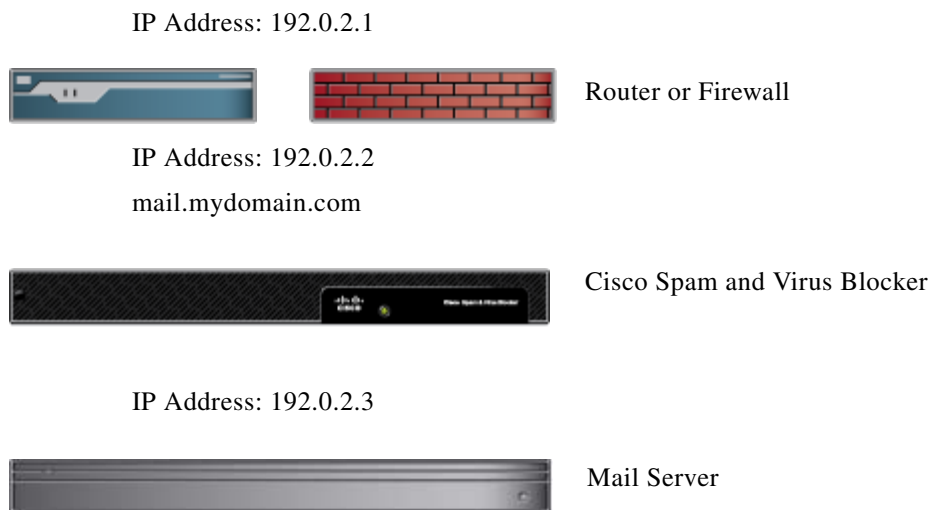
A DNS record is like an entry in an internet “phone book” for your domain. It translates a hostname, such as example.com, into an IP address. Included in the DNS record is an A record that maps the appliance hostname to its IP address and an MX record that directs incoming email to the correct mail server.

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 1-4](#)).

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

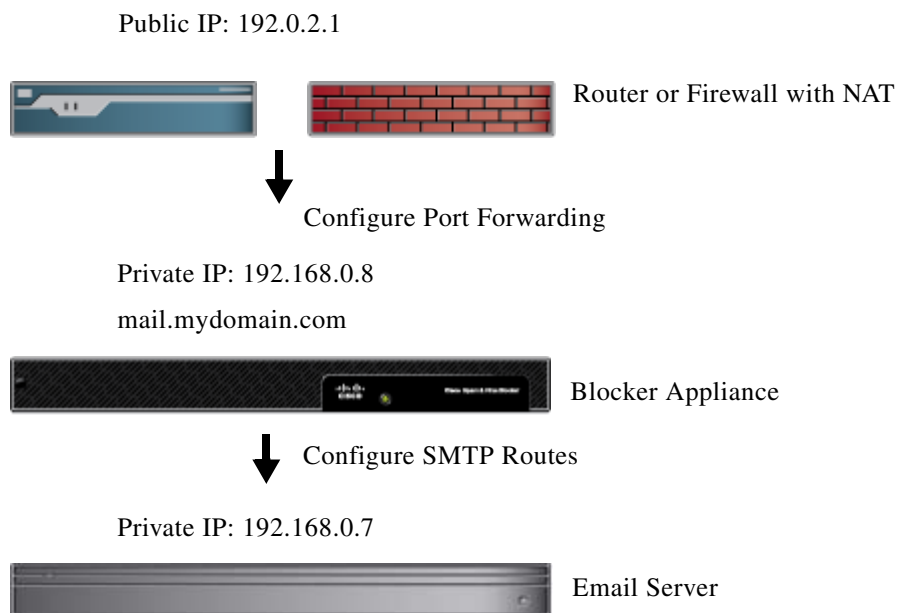
MX record: mydomain.com in MX mail.mydomain.com

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 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 Anti-Spam quarantine.
83	HTTPS	In	Blocker IPs	Used for viewing the Blocker Anti-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 “Preparing for Setup”, 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:

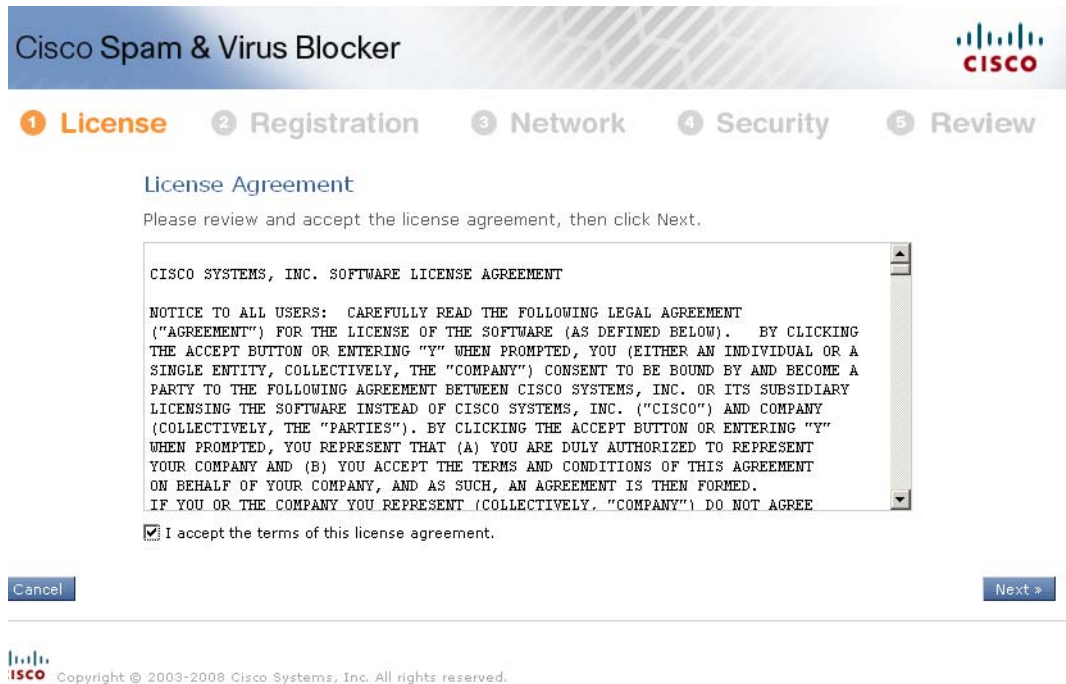
Access the System Setup Wizard from System Administration > System Setup Wizard

Figure 1-4 System Setup Wizard - Welcome Page



Step 1: Accept license agreement.

Figure 1-5 The System Setup Wizard - Accept License Page



Cisco Spam & Virus Blocker

1 License 2 Registration 3 Network 4 Security 5 Review

License Agreement


Please review and accept the license agreement, then click Next.

CISCO SYSTEMS, INC. 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 SYSTEMS, INC. OR ITS SUBSIDIARY LICENSING THE SOFTWARE INSTEAD OF CISCO SYSTEMS, INC. ("CISCO") 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

I accept the terms of this license agreement.

Cancel Next >

 Copyright © 2003-2008 Cisco Systems, Inc. All rights reserved.

Step 2: Enter registration information

Figure 1-6 The System Setup Wizard - Registration Page

1 License 2 **Registration** 3 Network 4 Security 5 Review

A 30 day evaluation license is active until you complete the product registration. After you register, Blocker is automatically updated with a permanent license.

Company Information

Company Name:

Address:
 (optional)

City: State/Province/Region:

Postal Code: Country:

Primary Contact

First Name: Last Name:

Email: Re-enter Email:

Business Phone: Mobile Phone: (optional)

Secondary Contact (optional)

First Name: Last Name:

Email: Re-enter Email:

Business Phone: Mobile Phone:

More Information

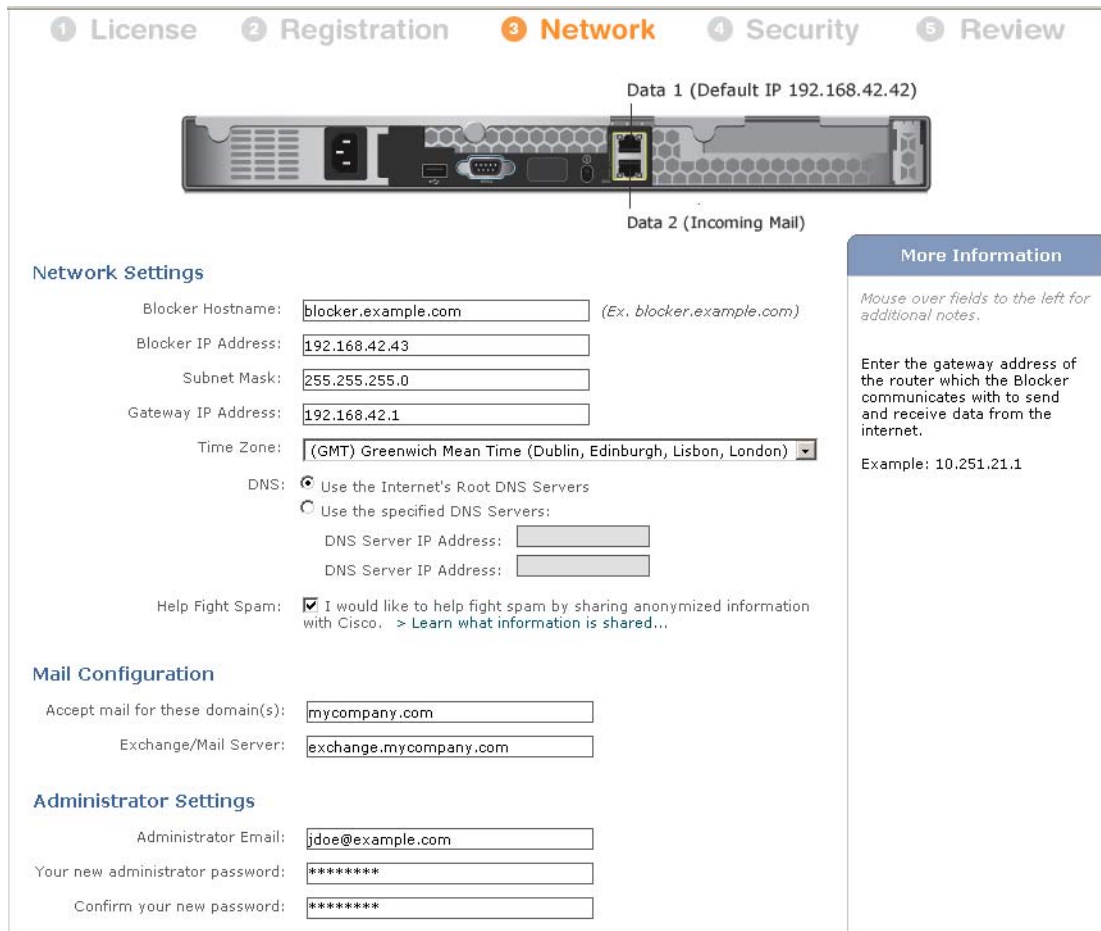
Mouse over fields to the left for additional notes.

Please enter your company's contact information and a primary point of contact so we may process your registration quickly. A valid email address is required for a permanent license.

Cancel Next >

Step 3: Enter network information

Figure 1-7 The System Setup Wizard - Network Page



1 License 2 Registration 3 Network 4 Security 5 Review

Data 1 (Default IP 192.168.42.42)

Data 2 (Incoming Mail)

Network Settings

Blocker Hostname: (Ex. blocker.example.com)

Blocker IP Address:

Subnet Mask:

Gateway IP Address:

Time Zone:

DNS: Use the Internet's Root DNS Servers
 Use the specified DNS Servers:
 DNS Server IP Address:
 DNS Server IP Address:

Help Fight Spam: I would like to help fight spam by sharing anonymized information with Cisco. > [Learn what information is shared...](#)

Mail Configuration

Accept mail for these domain(s):

Exchange/Mail Server:

Administrator Settings

Administrator Email:

Your new administrator password:

Confirm your new password:

More Information

Mouse over fields to the left for additional notes.

Enter the gateway address of the router which the Blocker communicates with to send and receive data from the internet.

Example: 10.251.21.1

Step 4: Set Anti-Spam and Anti-Virus Policy and install the configuration.**Figure 1-8 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:

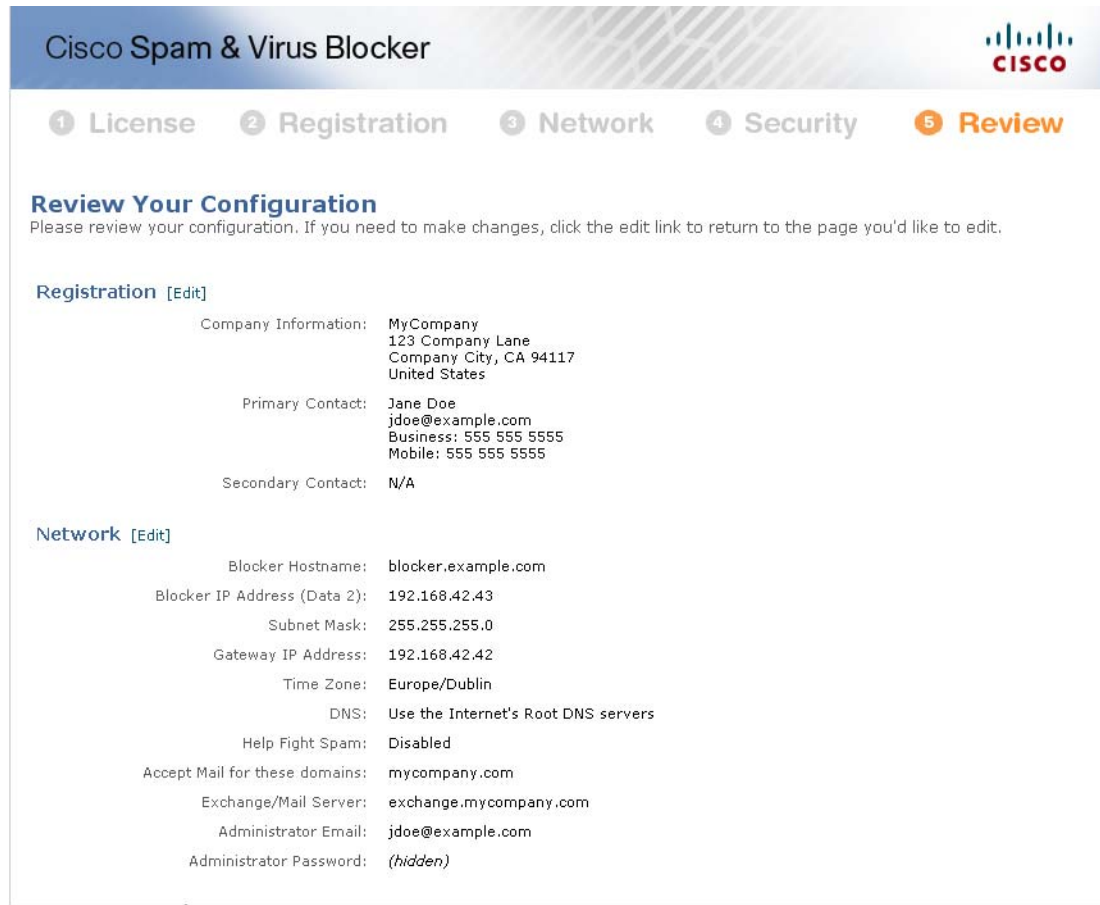
- Block Viruses

« Previous Cancel Next »

More Information

Mouse over fields to the left for additional notes.

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

Step 5: Review and Install the Configuration.**Figure 1-9 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)

Log back in to the appliance with the password you set in the Wizard. Once you log back into the appliance, the GUI opens to the “Next Steps” page which directs you to run a system test. Verify that your system is running. As a final test, send yourself an email from a personal email service, such as Gmail or Yahoo! email and view your Incoming Mail reports to ensure that the mail delivery is displayed.

Configuring an Active Directory Server Profile

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

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

After the Active Directory Wizard creates the LDAP server profile, use the System Administration > LDAP page to view the new profile and make additional changes.

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

To access the web-based Graphical User Interface (GUI) at the factory default address, open your web browser and point it to **http://192.168.42.42**.

The login screen is displayed:

Log in to 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.

All users accessing the GUI must log in. Type your username and password, and then click Login to access the GUI. You must use a supported web browser. You can log in with the admin account or with a specific user account you have created.

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, Mail Policies, Security Services, Network, and System Administration. The following chapters will describe each section, including the tasks you perform on pages within each section.

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

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

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

Monitor menu

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

Mail Policies menu

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

The Commit Changes Button

As you make configuration changes in the GUI, you must explicitly commit those changes by clicking the **Commit Changes** button. This button displays when you have uncommitted changes that need to be committed.

Figure 1-10 *Commit Changes Button*



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

Figure 1-11 *Commit Changes Page*

Uncommitted Changes

System Administration

System administration in general is handled primarily via the System Administration menu in the Graphical User Interface (GUI). Some system administration features are accessible only via the Command Line Interface (CLI).

Shutting Down the Blocker

To shut down your Blocker appliance, use the Shutdown / Reboot page available on the System Administration menu in the GUI.

**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.

Shutting down your appliance exits the Blocker, which allows you to safely power down the appliance. You may restart the appliance at a later time without losing any messages in the delivery queue. 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 reboot your Blocker use the Shutdown / Reboot page available on the System Administration menu in the GUI.

Rebooting your appliance restarts the Blocker, which allows you to safely power down and reboot the appliance. The default delay is thirty (30) seconds. The Blocker allows open connections to complete during the delay, after which it forcefully closes open connections. You may restart the appliance without losing any messages in the delivery queue.

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 1-12 **Open a Support Case**
Open a Technical Support Case

Technical Support Case		
Send Request to:	<input type="text"/> <i>Separate multiple email addresses with commas.</i>	
Contact Information:	Name: <input type="text"/>	
	Email: <input type="text"/>	
	Other Contact Information <i>(optional)</i>	
	Phone1: <input type="text"/>	
	Phone2: <input type="text"/> <i>(Mobile, Pager, etc.)</i>	
	Other: <input type="text"/>	
Issue Description:	Issue Subject: <input type="text"/>	
	Issue Description: <input type="text"/>	

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

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 AsyncOS 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
- AsyncOS 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
 - AsyncOS upgrades
 - SenderBase Network Participation sharing

Adding Users

The Blocker appliance provides two methods for adding user accounts: creating user accounts on the Blocker itself, and enabling user authentication using your own centralized authentication system, which can be an LDAP directory. You can manage users and connections to external authentication sources on the System Administration > Users page in the GUI.

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

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, Operator, Guest, Read-Only Operator, or Help Desk User. Each role contains differing levels of permissions within the system. After you have assigned a role, you specify a password for the user.

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.

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.)

Accessing the Appliance

You can access any IP interface you create on the appliance through a variety of services. By default, the following services are either enabled or disabled on each interface

Service	Default port	Enabled by default?	
		Management interface	New IP interfaces you create
FTP	21	No	No
Telnet	23	Yes	No
SSH	22	Yes	No
HTTP	80	Yes	No
HTTPS	443	Yes	No

Understanding the Email Pipeline

The following section provides an overview of how email is processed through the system, from reception to routing to delivery. Each feature is processed in order (from top to bottom) and is briefly described below. A full description of each feature can be found in the *AsyncOS for Email User Guide* and the *AsyncOS for Email Advanced User Guide*.

Shaded areas in the table represent processing that occurs in the Work Queue. You can test most of the configurations of features in this pipeline using the trace feature.

Email Pipeline for the Blocker Appliance: Receiving Email Features

Feature	Description
Host Access Table (HAT)	ACCEPT, REJECT, RELAY, or TCPREFUSE connections
Host DNS Sender Verification	Maximum outbound connections
Sender Groups	Maximum concurrent inbound connections per IP address
Envelope Sender Verification	Maximum message size and messages per connection
Sender Verification Exception Table	Maximum recipients per message and per hour
Mail Flow Policies	TCP listen queue size TLS: no/preferred/required SMTP AUTH: no/preferred/required Drop email with malformed FROM headers Always accept or reject mail from entries in the Sender Verification Exception Table. SenderBase on/off (IP profiling/flow control)
Received Header	Adds a received header to accepted email: on/off.
Default Domain	Adds default domain for “bare” user addresses.
Bounce Verification	Used to verify incoming bounce messages as legitimate.

Email Pipeline for the Blocker Appliance: Receiving Email Features

Domain Map	Rewrites the Envelope Recipient for each recipient in a message that matches a domain in the domain map table.
Recipient Access Table (RAT)	(Public listeners only) ACCEPT or REJECT recipients in RCPT TO plus Custom SMTP Response. Allow special recipients to bypass throttling.
Alias tables	Rewrites the Envelope Recipient. (Configured system-wide. aliasconfig is not a subcommand of listenerconfig.)
LDAP Recipient Acceptance	LDAP validation for recipient acceptance occurs within the SMTP conversation. If the recipient is not found in the LDAP directory, the message is dropped or bounced. LDAP validation can be configured to occur within the work queue instead.

Email Pipeline for the Blocker appliance: Routing and Delivery Features

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

Email Pipeline for the Blocker appliance: Routing and Delivery Features

Delivery limits	<ol style="list-style-type: none"> 1. Sets the default delivery interface. 2. Sets the total maximum number of outbound connections.
Domain-based Limits	Defines, per-domain: maximum outbound connections for each virtual gateway and for the entire system; the bounce profile to use; the TLS preference for delivery: no/preferred/required
Domain-based routing	Routes mail based on domain without rewriting Envelope Recipient.
Global unsubscribe	Drops recipients according to specific list (configured system-wide).
Bounce profiles	Undeliverable message handling. Configurable per listener, per Destination Controls entry, and via message filters.

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

** Can send messages to the Spam Quarantine.



CHAPTER 2

Accepting Mail

Revised: January 22, 2008, OL-18046-02

Accepting Mail: Overview

After you have configured the basic settings of your Blocker appliance via the System Setup Wizard, you are now ready to begin tailoring the configuration of your Blocker to receive email. This chapter discusses some of the key features available to you as you begin to configure listeners on the appliance to handle receiving email. 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.

The concept of the Host Access Table (HAT) is introduced. The Host Access Tables (HATs) of public listeners — with their specific sender groups and mail flow policies — provide the underlying framework that makes possible the Mail Flow Monitor feature.

Configuring the Blocker to Accept Mail (Public Listeners)

The Blocker's operating system allows it to function as the inbound email gateway for your enterprise, servicing SMTP connections from the Internet, accepting messages, and relaying messages to the appropriate systems.

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

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

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

Use the Listeners page (Network > Listeners) to configure listeners that run over available IP interfaces on the Blocker appliance.

Public and Private Listeners

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

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. For email delivery and Virtual Gateways, each IP interface acts as one Virtual Gateway address with a specific IP address and hostname. For more information, see the “Advanced Networking” chapter in the *AsyncOS for Email Advanced User Guide*.

Incoming Relays

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), filtering servers, etc. at the edge of the network. In this type of configuration, the IP address of the external machine is not automatically known by the Blocker. Instead, mail appears to originate from the local MX/MTA (the incoming relay) rather than from the external machine. IronPort Anti-Spam depends on accurate IP addresses for external senders so it is vital for the Blocker appliance to have this information.

You should only enable this feature if you have a local MX/MTA relaying mail to your Blocker appliance.

When configuring an incoming relay, you specify the names and IP addresses of all of the internal MX/MTAs connecting to the Blocker appliance, as well as the header used to store the originating IP address. You have two options for specifying the header: a custom header or an existing received header.

For details about incoming relays, see the “Anti-Spam” chapter in the *AsyncOS for Email User Guide*.

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 system, and one good way of determining how to filter mail is to determine who sent the mail to you. For example, when sorting through the mail that comes to you via the post office, you might keep the letters sent from your grandmother, but throw away letters from Jim’s Spandalot Outlet, simply by noting who sent you the mail. Similarly, your Blocker can perform actions on your email based on 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 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 an appliance has properties that you can configure to modify the behavior of the message it receives. One of the first configurable features that influences a listener’s behavior is its Host Access Table (HAT).

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

Remote Host Definition	Rule
------------------------	------

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

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

Extending the basic syntax, HATs in 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
Remote Host	
Remote Host	
Remote Host	
...	

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

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

Default HAT Entries

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

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

Mail Flow Policies: Access Rules and Parameters

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

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

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. Domain Keys signing is available only on RELAY mail flow policies.

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. The CONTINUE rule is used to facilitate the editing of the HAT in the Graphical User Interface (GUI).

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.
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.

Parameter	Description
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 IronPort 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. For more information about HAT significant bits, see the “HAT Significant Bits Feature” in chapter 1 of the <i>AsyncOS for Email Advanced User Guide</i> .
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). For more information on configuring DHAP for LDAP accept queries, see “LDAP Queries” in the <i>AsyncOS for Email Advanced User Guide</i> .
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	

Parameter	Description
Allow TLS Connections	Deny, Prefer, or Require Transport Layer Security (TLS) in SMTP conversations for this listener.
SMTP Authentication	Allows, disallow, or requires SMTP Authentication from remote hosts connecting to the listener. SMTP Authentication is described in detail in the “LDAP Queries” chapter of the <i>AsyncOS for Email Advanced User Guide</i> .
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. For more information, see the <i>AsyncOS for Email Advanced User Guide</i> .
Conformance Level	Set the SPF/SIDF conformance level. You can choose from SPF, SIDF or SIDF Compatible. For details, see <i>AsyncOS for Email Advanced User Guide</i> .
Downgrade PRA verification result if 'Resent-Sender:' or 'Resent-From:' were used:	If you choose a conformance level of SIDF compatible, configure whether you want to downgrade Pass result of the PRA Identity verification to None if there are Resent-Sender: or Resent-From: headers present in the message. You may choose this option for security purposes.
HELO Test	Configure whether you want to perform a test against the HELO identity (Use this for SPF and SIDF Compatible conformance levels).
Untagged Bounces	
Consider Untagged Bounces to be Valid	Applies only if bounce verification tagging (discussed in detail in the <i>AsyncOS for Email Advanced User Guide</i>) 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	
	See Sender Verification in the <i>AsyncOS for Email User Guide</i> .
Exception Table	
Use Exception Table	Use the sender verification domain exception table. You can only have one exception table, but you can enable it per mail flow policy. See the <i>AsyncOS for Email User Guide</i> for more information.

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

HAT Variable Syntax

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

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

Using HAT Variables

These variables can be used with the `smtp_banner_text` and `max_rcpts_per_hour_text` advanced HAT parameters shown in the “Customizing Listeners” chapter in the *AsyncOS for Email Advanced User Guide*.

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

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)

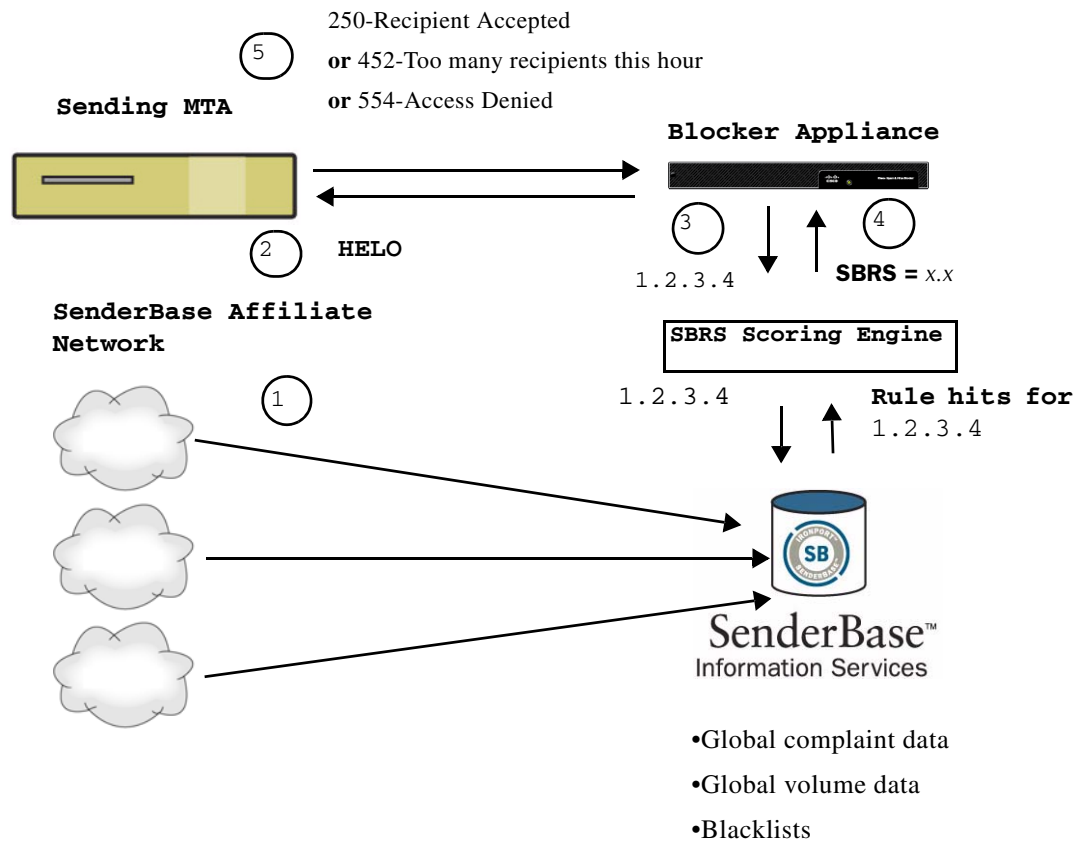
The SenderBase Reputation Service (SBRS) score is a numeric value assigned to an IP address based on information from the SenderBase Reputation Service. The SenderBase Reputation Service aggregates data from over 25 public blacklists and open proxy lists, and combines this data with global data from 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. For more information, refer to “SenderBase Reputation Rule” and “Bypass Anti-Spam System Action” in the “Using Message Filters to Enforce Email Policies” chapter in the *AsyncOS for Email Advanced User Guide*.)

Figure 2-1 The Senderbase Reputation Service



1. SenderBase affiliates send real-time, global data
2. Sending MTA opens connection with the Blocker
3. Blocker appliance checks global data for the connecting IP address
4. SenderBase Reputation Service calculates the probability this message is spam and assigns a SenderBase Reputations Score
5. 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.

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



CHAPTER 3

Mail Routing For Users and Groups

Revised: January 22, 2008, OL-18046-02

Creating Mail Rules for Specific Users , Groups, or Domains

The Blocker allows you a lot of flexibility in defining the actions to take for mail recipients, and it provides several options for defining groups of users. When you define the users for whom you will receive mail, you can apply these rules to a 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 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.



Note

When you run the System Setup Wizard, you create a listener for incoming mail by default. You can implement the following settings on your default listener, or you can create more listeners. For example, you might want to create a separate listener for outgoing mail.

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 a New RAT Entry, page 3-37](#)

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

Adding a New RAT Entry

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 checked="" type="checkbox"/> Bypass LDAP Accept Queries for this Recipient
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 Recieve 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 Recieve 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.

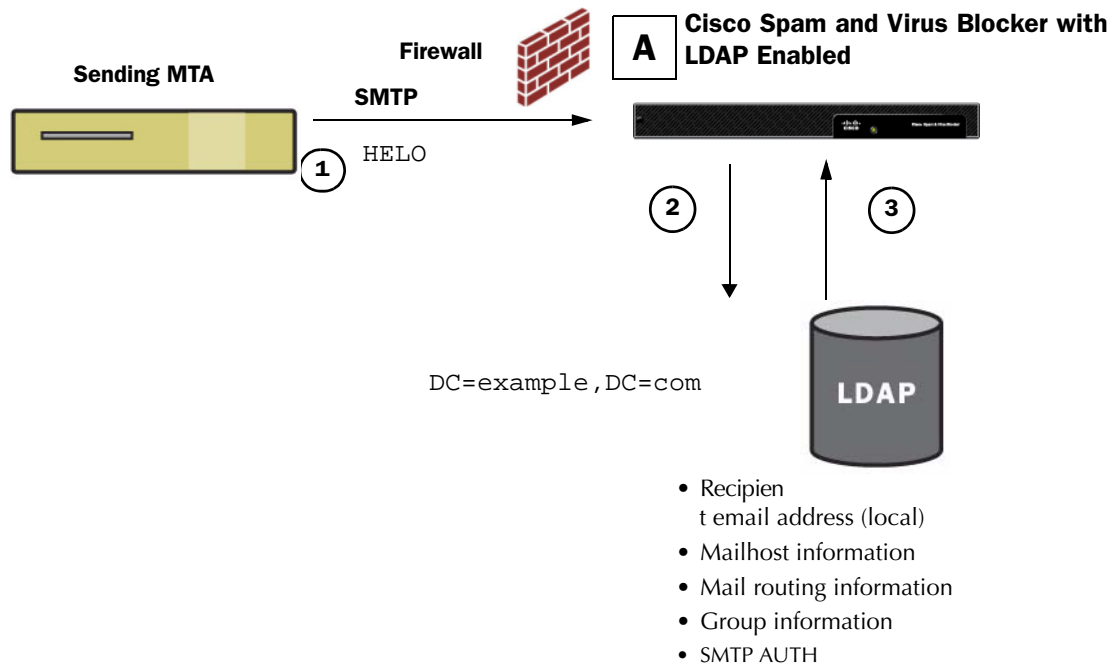
For details about deleting, importing or exporting RAT entries, see “Configuring the Gateway to Receive Email” in the *AsyncOS for Email User Guide*.

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 an Active Directory Server Profile, page 1-12](#).

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 [Content Filters Overview, page 6-63](#).

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 (formerly known as “iPlanet 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.
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: January 22, 2008, OL-18046-02

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 *sendergroup* 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 2-25](#).

Basic Steps for Configuring Outbound Mail

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

The screenshot shows the 'Sender Group Settings' for the 'RELAYLIST' group. The settings are as follows:

Field	Value
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

Buttons: '<< Back to HAT Overview' and 'Edit Settings...'

Below this is the 'Find Senders' section with a search input field and a 'Find' button.

Below that is the 'Sender List: Display All Items in List' section with an 'Add Sender...' button and the message '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

The screenshot shows the 'Sender Groups' interface for the 'OutgoingMail' listener. It displays a table with the following data:

Order	Sender Group	SenderBase™ Reputation Score	Mail Flow Policy	Delete
1	RELAYLIST		RELAYED	
	ALL		BLOCKED	

Buttons: 'Add Sender Group...', 'Import HAT...', 'Edit Order...', and '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 4-49](#).
- using the content filter action, Add Disclaimer Text. For details, see the [Content Filter Actions, page 6-69](#).

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

You must create the disclaimer before you can add it to the listener. To create a text resource, go the Mail Policies > Text Resources and click **Add text resource**. Select “Disclaimer” as Type, and enter a name and text for the disclaimer. For more details, see “Adding Text Resources” in the *AsyncOS for Email User Guide*.

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: January 22, 2008, OL-18046-02

Virus and Spam Protection Overview

Virus Scanning

The Blocker includes integrated virus scanning engines from Sophos, Plc. You can configure the appliance to scan messages for viruses (based on the matching incoming or outgoing mail policy), and, if a virus is found, to perform different actions on the message (including “repairing” the message of viruses, modifying the subject header, adding an additional X-header, sending the message to an alternate address or mailhost, archiving the message, or deleting the message).

Anti-spam Scanning

The Blocker offers a unique, layered approach to stopping spam at the email gateway. The first layer of spam control, reputation filtering, allows you to classify email senders and restrict access to your email infrastructure based on senders’ trustworthiness as determined by the IronPort SenderBase™ Reputation Service. The second layer of defense, scanning, is powered by IronPort Anti-Spam™ technology. Coupled together, reputation filtering and anti-spam scanning offer the most effective and highest performing anti-spam solution available today.

Using the 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.

The unique, two-layer approach to fighting spam with the Blocker provides you with a powerful and unprecedented flexibility to manage and protect your enterprise email gateway.

Enabling Virus Protection

The Blocker includes integrated virus scanning engines from Sophos, Plc. You can configure the appliance to scan messages for viruses (based on the matching incoming or outgoing mail policy), and, if a virus is found, to perform different actions on the message (including “repairing” the message of viruses, modifying the subject header, adding an additional X-header, sending the message to an alternate address or mailhost, archiving the message, or deleting the message).

If enabled, virus scanning is performed in the “work queue” on the appliance, immediately after Anti-Spam scanning.

Steps to Enable Virus Protection

You can enable a virus scanning engine when you run the System Setup Wizard. Or, you can enable and modify the virus scanning engine global configuration settings via Security Services > Sophos 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
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 per-user anti-virus settings for a mail policy is essentially the same for incoming or outgoing mail.

Individual policies (not the default) have an additional field to “Use Default” settings. Select this setting to inherit the default mail policy settings.

You enable anti-virus actions on a per-recipient basis using the Email Security Feature: the Mail Policies > Incoming or Outgoing Mail Policies pages. 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. 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.

Figure 5-1 Anti-Virus Settings for a Mail Policy (not default) 1 of 2

No Changes Pending

Mail Policies: Anti-Virus

Anti-Virus Settings	
Policy:	DEFAULT
Enable Anti-Virus Scanning for This Policy:	<input checked="" type="radio"/> Yes <input checked="" type="checkbox"/> Use Sophos Anti-Virus <input type="radio"/> No
Message Scanning	
	Scan for Viruses only <input type="text" value="Scan for Viruses only"/>
	<input type="checkbox"/> Drop infected attachments if a virus is found <input checked="" type="checkbox"/> (recommended) Include an X-header with the Anti-Virus scanning results in messages
Repaired Messages:	
Action Applied to Message:	<input type="text" value="Deliver As Is"/>
Archive Original Message:	<input checked="" type="radio"/> No <input type="radio"/> Yes
Modify Message Subject:	<input checked="" type="radio"/> No <input type="radio"/> Prepend <input type="radio"/> Append
	<input type="text" value="[WARNING: VIRUS REMOVED]"/>
Advanced	Optional settings for custom header and message delivery.

Figure 5-2 Anti-Virus Settings for a Mail Policy (not default) 2 of 2

Encrypted Messages:	
Action Applied to Message:	Deliver As Is
Archive Original Message:	<input checked="" type="radio"/> No <input type="radio"/> Yes
Modify Message Subject:	<input checked="" type="radio"/> No <input type="radio"/> Prepend <input type="radio"/> Append
	[WARNING : MESSAGE ENCRYPT]
▶ Advanced	Optional settings for custom header and message delivery.
Unscannable Messages:	
Action Applied to Message:	Deliver As Is
Archive Original Message:	<input type="radio"/> No <input checked="" type="radio"/> Yes
Modify Message Subject:	<input checked="" type="radio"/> No <input type="radio"/> Prepend <input type="radio"/> Append
	[WARNING : A/V UNSCANNABLE]
▶ Advanced	Optional settings for custom header and message delivery.
Virus Infected Messages:	
Action Applied to Message:	Drop Message
Archive Original Message:	<input type="radio"/> No <input checked="" type="radio"/> Yes
Modify Message Subject:	<input type="radio"/> No <input checked="" type="radio"/> Prepend <input type="radio"/> Append
	[WARNING : VIRUS DETECTED]
▶ Advanced	Optional settings for custom header and message delivery.

Enabling Spam Protection

The Blocker offers a unique, layered approach to stopping spam at the email gateway. The first layer of spam control, reputation filtering, allows you to classify email senders and restrict access to your email infrastructure based on senders’ trustworthiness as determined by the 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, outside of 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 an 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 5-3 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 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.

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, 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 Installation 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:

```
sibl<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 antis spam 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*.

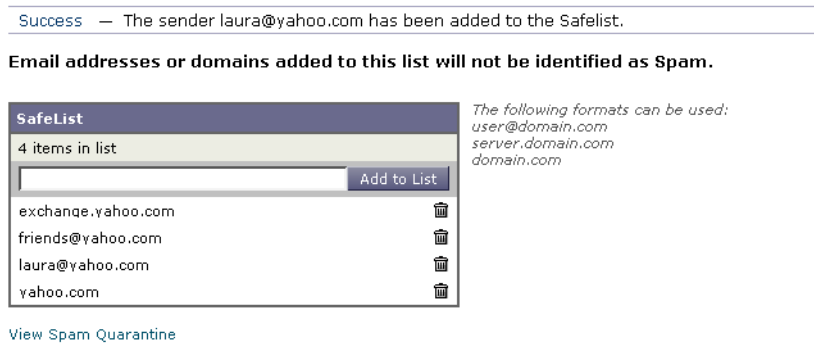
Adding Entries to Safelists

End users can add senders to safelists in two ways:

Method 1:

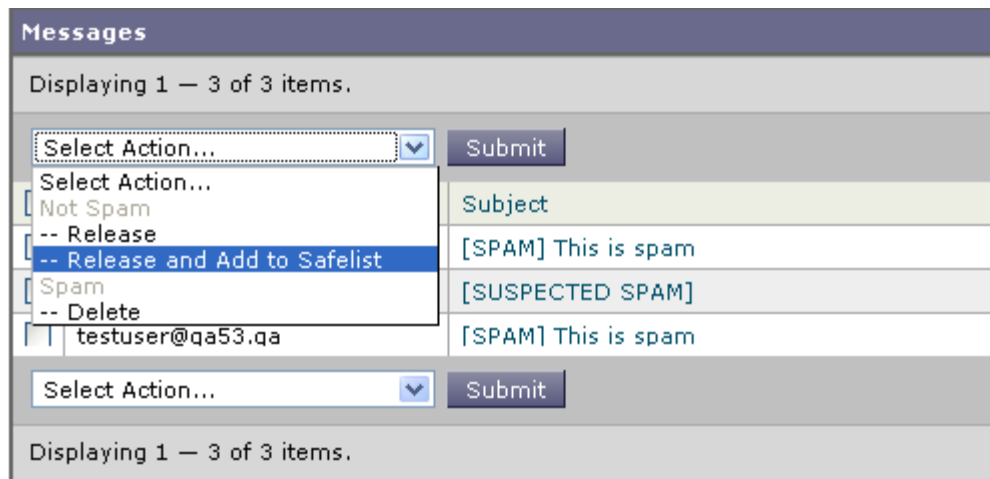
1. From the Spam Quarantine, select the Options drop-down menu.
2. Choose Safelist.
3. From the Safelist dialog box, enter the email address or domain.
4. Click Add to List.

Figure 5-4 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 5-5 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

Enforcing Policies Using Content Filters

Revised: January 22, 2008, OL-18046-02

Content Filters Overview

Email Security Manager policies allow you to create content filters to be applied to messages on a per-recipient or per-sender basis. Content filters are 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 message filters processing and anti-spam and anti-virus scanning have been performed on a message.

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.

To Create a Content Filter

To create the content filter:

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 “Content Filter Conditions” on page 64.

6. Select a condition to add.

7. Click **Add Action**.

For information about the actions you can apply, see “Content Filter Actions” on page 69.

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.

Content Filter Conditions

Specifying conditions in content filters is optional.

In the content filter conditions, when you add filter rules that search for patterns in the message body or attachments, you can specify the minimum threshold for the number of times the pattern must be found. When 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”).

The following shows the content filter condition screen:

Figure 6-1 Content Filter Conditions

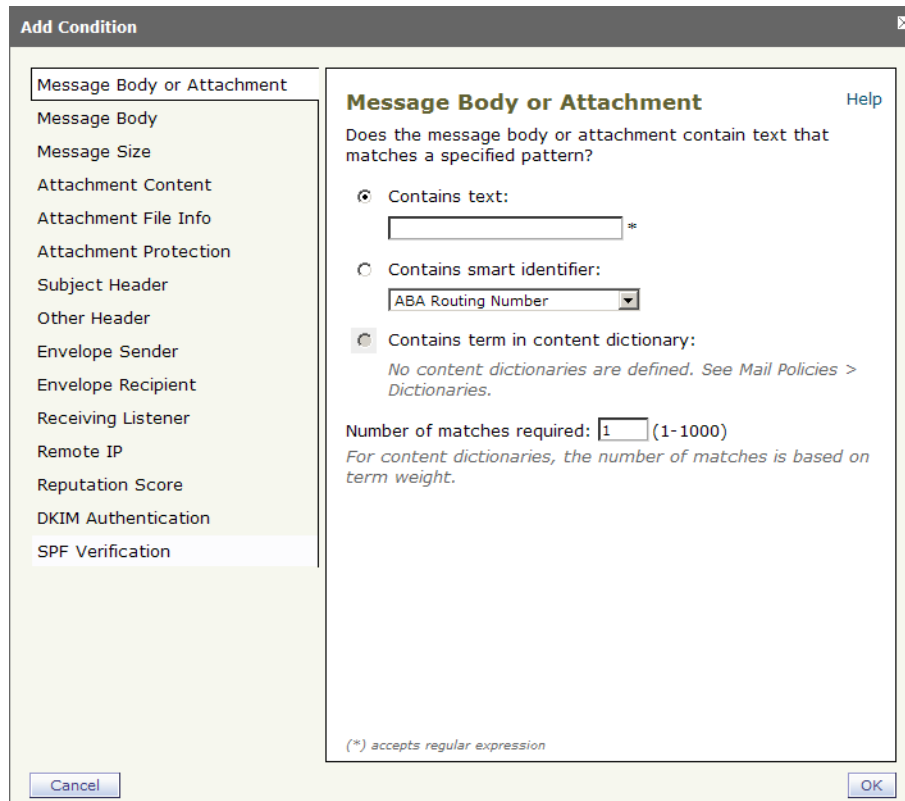


Table 6-1 Content Filter Conditions

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 <dictionary name>?</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. The IP address pattern is specified using the allowed hosts notation described in “Sender Group Syntax” in the <i>AsyncOS for Email User Guide</i> , except for the SBO, SBRS, dnslist notations and the special keyword ALL.
Reputation Score	What is the sender’s SenderBase Reputation Score? The Reputation Score rule checks the SenderBase Reputation Score against another value.
DKIM Authentication	Did DKIM authentication pass, partially verify, return temporarily unverifiable, permanently fail, or were no DKIM results returned?
SPF Verification	What was the SPF verification status? This filter rule allows you to query for different SPF verification results. For more information about SPF, see <i>AsyncOS for Email Advanced User Guide</i> .

Content Filter Actions

At least one action must be defined for each content filter.

Actions are performed in order on messages, so consider the order of actions when defining multiple actions for a content filter.

When you configure a quarantine action for messages that match Attachment Content conditions, Message Body or Attachment conditions, Message body conditions, or the Attachment content conditions, you can view the matched content in the quarantined message. When you display the message body, the matched content is highlighted in yellow. You can also use the `$MatchedContent` action variable to include the matched content in the message subject. For more information, see 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

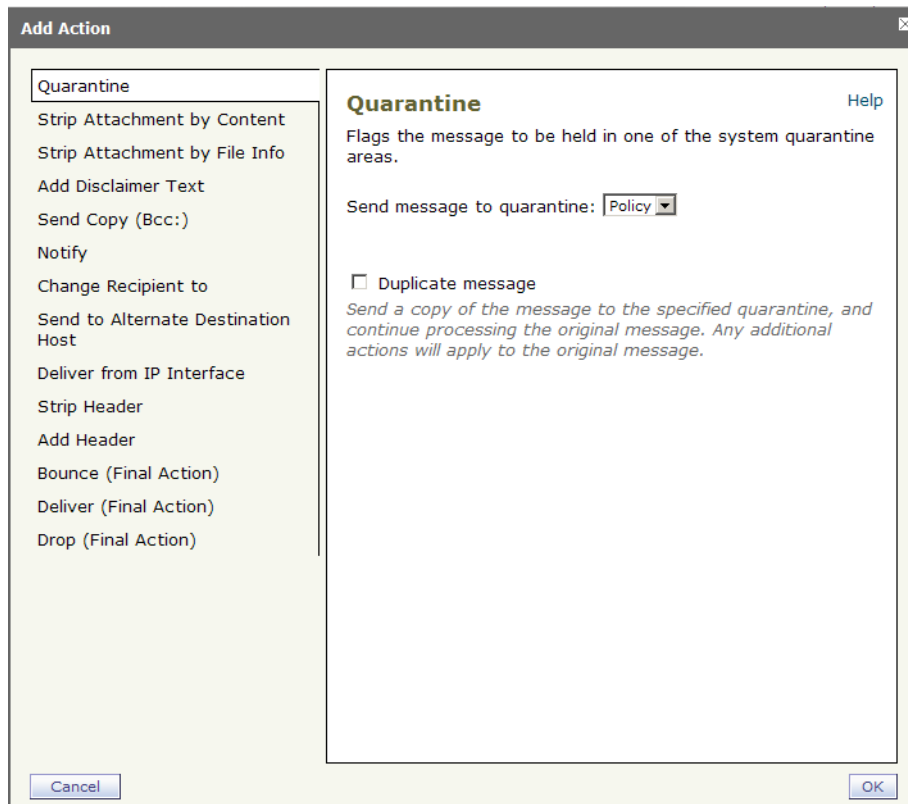


Table 6-2 Content Filter Actions

Action	Description
Quarantine	<p>Quarantine. Flags the message to be held in one of the system quarantine areas. See the “Filters” chapter in the <i>AsyncOS for Email Advanced User Guide</i> for more information.</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 already created disclaimer text in order to use this content filter action.</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>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.



CHAPTER 7

Monitoring Your Email

Revised: January 22, 2008, OL-18046-02

Email Monitoring Overview

You can use message tracking queries or the reporting tool to help troubleshoot email problems and monitor email that is entering your network.

You can use the message tracking tool to discover the status of a particular email, and you can use the Email Monitor to get an understanding of the email that is entering your network. For example, you can see what percentage of your email is spam, who is sending you the most spam email,

Finding Emails Using Message Tracking

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.

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.

Using the Email Security Monitor

You use Email Security Monitor to monitor domains that send mail to your company. 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.

Searching and Email Security Monitor

Many of the Email Security Monitor pages include a search form. You can search for different types of items:

- IP Address
- domain
- network owner
- internal users
- destination domain
- internal sender domain
- internal sender IP address
- outgoing domain deliver status

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.

Email Security Monitor Overview Page

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.
- See current virus outbreak information based on information available at the Threat Operations Center (TOC).

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
- Resource Conservation
- Delivery Suspended
- Receiving Suspended
- Work Queue Paused
- Offline

See the Managing and Monitoring chapter in the *AsyncOS for Email Advanced User Guide* for more information.

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.

Categorizing Email

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.
- **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.
- **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.

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.

Outgoing Destinations

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?

Outgoing Senders

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?

The Delivery Status Page

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.

Internal User Details

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.

The Content Filters Page

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.

Virus Types Page

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.

TLS Connections Page

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?

The System Capacity Page

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.

The System Status Page

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 AsyncOS operating system installed.
- The installation date of the AsyncOS 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.

Debugging Mail Flow Using Test Messages: Trace

Yet 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. The test message is not actually sent. The Trace page can be a powerful troubleshooting or debugging tool, especially if you have combined many of the advanced features available on the Blocker appliance.



CHAPTER 8

Advanced Topics

Revised: January 22, 2008, OL-18046-02

Advanced Topics Overview

The Blocker is a robust appliance with a variety of advanced options that are not discussed in this guide. These advanced topics are discussed in one of the following reference guides:

- *AsyncOS for Email Advanced User Guide*
- *AsyncOS for Email User Guide*

PDF versions of these guides are included in the documentation CD shipped with your Blocker appliance.



Note

Some of the options discussed in the guides for Email Security appliances are not available for the Blocker appliance. This includes IronPort email encryption, centralized reporting, centralized management, McAfee virus protection, image analysis, external quarantines, remote access, and external user authentication.

Listing of Advanced Topics

The following section includes a brief description of some advanced topics and their locations in the AsyncOS for Email user guides, including chapters and book names.

- **LDAP Configuration.** This guide outlines the steps for creating an LDAP profile. For detailed information about configuring LDAP profiles, see “LDAP Queries” in *AsyncOS for Email Advanced User 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. SMTP Routes are discussed in the *AsyncOS for Email Advanced User 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. Domain Keys and DKIM signing are covered in *AsyncOS for Email Advanced User 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. Details about logging are covered in the *AsyncOS for Email Advanced User 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. Details about how to create a bounce profile are in the *AsyncOS for Email Advanced User 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).
- **SMTP Authentication.** Blocker provides support for SMTP authentication. SMTP Auth is a mechanism for authenticating clients connected to an SMTP server.

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. In addition, this release implements a read-only subset of MIB-II as defined in RFCs 1213 and 1907. (For more information on SNMP, see RFCs 1065, 1066, and 1067.) For more information, see *AsyncOS for Email Advanced User Guide*.
- **Customizing Listeners.** The Blocker supports many options for customizing listeners, such as encrypting SMTP conversations using TLS. For details, see “Customizing Listeners” in the *AsyncOS for Email Advanced User Guide*.



GLOSSARY

Revised: January 22, 2008,
OL-18046-02

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 malware allow anti-virus scanners to detect and remove not only viruses, but also all types of malicious software.

B

- Blacklist** A list of known bad senders. By default, senders in the Blacklist sender group of a public listener are rejected by the parameters set in the \$BLOCKED mail flow policy.

C

- Character Set (Double-byte)** Double Byte Character Sets are foreign-language character sets requiring more than one byte of information to express each character.
- CIDR Notation** Classless Inter-Domain Routing. A convenient shorthand for describing a range of IP addresses within their network contexts using an arbitrary number of bits. Using this notation, you note the network prefix part of an address by adding a forward slash (/) followed by the number of bits used for the network part. Thus a Class C network can be described in prefix notation as 192.168.0.1/24. A CIDR specification of 206.13.1.48/25 would include any address in which the first 25 bits of the address matched the first 25 bits of 206.13.1.48..
- Content Filters** Content-based filters used to process messages during the Per-Recipient Scanning phase of the work queue in the email pipeline. Content filters are evoked after Message filters, and act on individual splintered messages.
- 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	<p>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 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.</p> <p>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.</p>
Log Subscription	<p>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.</p>
M	
Mail Flow Policies	<p>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.</p>
MAIL FROM	<p>See Envelope Sender.</p>
Maximum Number of Retries	<p>The maximum number of times that redelivery of a soft bounced message will be attempted before being hard bounced.</p>
Maximum Time in Queue	<p>The maximum length of time that a soft bounced message will stay in the email queue for delivery before being hard bounced.</p>
MTA	<p>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.</p>
MUA	<p>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.</p>
MX Record	<p>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.</p>

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: January 22, 2008, OL-18046-02

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-===}

