



Cisco AsyncOS 9.1 for Email User Guide

April 2, 2015

Cisco Systems, Inc.

www.cisco.com

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco website at www.cisco.com/go/offices.

<https://t.me/learningnets>

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 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.

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned 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. (1110R)

Cisco AsyncOS 9.1 for Email User Guide

© 2015 Cisco Systems, Inc. All rights reserved.



CHAPTER 1**Getting Started with the Cisco Email Security Appliance 1-1**

- What's New in This Release 1-1
 - What's New in Cisco AsyncOS 9.1 for Email 1-2
 - What's New in Cisco AsyncOS 9.0 for Email 1-3
- Where to Find More Information 1-7
 - Documentation 1-7
 - Training 1-8
 - Cisco Notification Service 1-8
 - Knowledge Base 1-9
 - Cisco Support Community 1-9
 - Cisco Customer Support 1-9
 - Third Party Contributors 1-9
 - Cisco Welcomes Your Comments 1-10
 - Registering for a Cisco Account 1-10
- Cisco Email Security Appliance Overview 1-10
 - Supported Languages 1-11

CHAPTER 2**Accessing the Appliance 2-1**

- Web-based Graphical User Interface (GUI) 2-1
 - Browser Requirements 2-1
 - Accessing the GUI 2-1
- Changing Configuration Settings 2-3
 - Configuration Changes 2-3
 - Commit or Abandoning Changes 2-3
- Command Line Interface (CLI) 2-3
 - Command Line Interface Conventions 2-3
 - General Purpose CLI Commands 2-7

CHAPTER 3**Setup and Installation 3-1**

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

Installation Scenarios	3-3
Physically Connecting the Email Security Appliance to the Network	3-5
Configuration Scenarios	3-5
Preparing for System Setup	3-8
Determine Method for Connecting to the Appliance	3-8
Determining Network and IP Address Assignments	3-9
Gathering the Setup Information	3-11
Using the System Setup Wizard	3-13
Accessing the Web-Based Graphical User Interface (GUI)	3-14
Defining Basic Configuration Using the Web-Based System Setup Wizard	3-14
Setting up the Connection to Active Directory	3-22
Proceeding to the Next Steps	3-23
Accessing the Command Line Interface (CLI)	3-23
Running the Command Line Interface (CLI) System Setup Wizard	3-24
Configuring your system as an Enterprise Gateway	3-37
Verifying Your Configuration and Next Steps	3-37

CHAPTER 4

Understanding the Email Pipeline	4-1
Overview of the Email Pipeline	4-1
Email Pipeline Flows	4-1
Incoming / Receiving	4-4
Host Access Table (HAT), Sender Groups, and Mail Flow Policies	4-5
Received: Header	4-5
Default Domain	4-5
Bounce Verification	4-5
Domain Map	4-6
Recipient Access Table (RAT)	4-6
Alias Tables	4-6
LDAP Recipient Acceptance	4-6
SMTP Call-Ahead Recipient Validation	4-6
Work Queue / Routing	4-7
Email Pipeline and Security Services	4-7
LDAP Recipient Acceptance	4-7
Masquerading or LDAP Masquerading	4-8
LDAP Routing	4-8
Message Filters	4-8
Email Security Manager (Per-Recipient Scanning)	4-8
Quarantines	4-10
Delivery	4-10

Virtual gateways	4-11
Delivery Limits	4-11
Domain-Based Limits	4-11
Domain-Based Routing	4-11
Global Unsubscribe	4-11
Bounce Limits	4-12

CHAPTER 5**Configuring the Gateway to Receive Email 5-1**

Overview of Configuring the Gateway to Receive Email	5-1
Working with Listeners	5-2
Configuring Global Settings for Listeners	5-6
Settings for Messages Containing Multiple Encodings: localeconfig	5-8
Listening for Connection Requests by Creating a Listener via the GUI	5-8
Partial Domains, Default Domains, and Malformed MAIL FROMs	5-12
Listening for Connection Requests by Creating a Listener via the CLI	5-13
Advanced HAT Parameters	5-14
Enterprise Gateway Configuration	5-15

CHAPTER 6**Sender Reputation Filtering 6-1**

Overview of Sender Reputation Filtering	6-1
SenderBase Reputation Service	6-1
SenderBase Reputation Score (SBRS)	6-2
How SenderBase Reputation Filters Work	6-3
Recommended Settings for Different Sender Reputation Filtering Approaches	6-4
Editing Sender Reputation Filtering Score Thresholds for a Listener	6-5
Testing Sender Reputation Filtering Using the SBRS	6-6
Monitoring the Status of the SenderBase Reputation Service	6-7
Entering Low SBRS Scores in the Message Subject	6-7

CHAPTER 7**Defining Which Hosts Are Allowed to Connect Using the Host Access Table (HAT) 7-1**

Overview of Defining Which Hosts Are Allowed to Connect	7-1
Default HAT Entries	7-2
Defining Remote Hosts into Sender Groups	7-3
Sender Group Syntax	7-4
Sender Groups Defined by Network Owners, Domains, and IP Addresses	7-4
Defining Sender Groups by SenderBase Reputation Score	7-6
Sender Groups Defined by Querying DNS Lists	7-7
Defining Access Rules for Email Senders Using Mail Flow Policies	7-8

HAT Variable Syntax	7-9
Understanding Predefined Sender Groups and Mail Flow Policies	7-11
Handling Messages from a Group of Senders in the Same Manner	7-13
Creating a Sender Group for Message Handling	7-13
Adding a Sender to an Existing Sender Group	7-14
Rearranging the Order of the Rules to Perform for Incoming Connections	7-14
Searching for Senders	7-15
Defining Rules for Incoming Messages Using a Mail Flow Policy	7-15
Defining Default Values for Mail Flow Policies	7-20
Working with the Host Access Table Configuration	7-21
Exporting the Host Access Table Configuration to an External File	7-21
Importing the Host Access Table Configuration from an External File	7-21
Using a List of Sender Addresses for Incoming Connection Rules	7-22
SenderBase Settings and Mail Flow Policies	7-23
Timeouts for SenderBase Queries	7-23
HAT Significant Bits Feature	7-24
Verifying Senders	7-28
Sender Verification: Host	7-28
Sender Verification: Envelope Sender	7-29
Implementing Sender Verification — Example Settings	7-31
Testing Your Settings for Messages from Unverified Senders	7-36
Sender Verification and Logging	7-37
Enabling Host DNS Verification via the CLI	7-38

CHAPTER 8

Accepting or Rejecting Connections Based on Domain Name or Recipient Address 8-1

Overview of Accepting or Rejecting Connections Based on the Recipient's Address	8-1
Overview of the Recipient Access Table (RAT)	8-2
Accessing the RAT	8-2
Editing the Default RAT Entry	8-2
Domains and Users	8-3
Adding Domains and Users For Which to Accept Messages	8-3
Rearranging the Order of Domains and Users in the Recipient Access Table	8-5
Exporting the Recipient Access Table to an External File	8-5
Importing the Recipient Access Table from an External File	8-6

CHAPTER 9

Using Message Filters to Enforce Email Policies 9-1

Overview	9-1
Components of a Message Filter	9-2

Message Filter Rules	9-2
Message Filter Actions	9-2
Message Filter Example Syntax	9-3
Message Filter Processing	9-4
Message Filter Order	9-4
Message Header Rules and Evaluation	9-5
Message Bodies vs. Message Attachments	9-5
Thresholds for Matches in Content Scanning	9-6
AND Test and OR Tests in Message Filters	9-9
Message Filter Rules	9-10
Filter Rules Summary Table	9-10
Regular Expressions in Rules	9-17
Smart Identifiers	9-21
Description and Examples of Message Filter Rules	9-22
Message Filter Actions	9-49
Filter Actions Summary Table	9-49
Action Variables	9-56
Matched Content Visibility	9-58
Description and Examples of Message Filter Actions	9-58
Attachment Scanning	9-78
Message Filters for Scanning Attachments	9-78
Image Analysis	9-80
Configuring the Image Analysis Scanning Engine	9-80
Configuring the Message Filter to Perform Actions Based on Image Analysis Results	9-83
Notifications	9-85
Examples of Attachment Scanning Message Filters	9-86
Using the CLI to Manage Message Filters	9-89
Creating a New Message Filter	9-90
Deleting a Message Filter	9-91
Moving a Message Filter	9-91
Activating and Deactivating a Message Filter	9-91
Importing Message Filters	9-95
Exporting Message Filters	9-96
Viewing Non-ASCII Character Sets	9-96
Displaying a Message Filter List	9-96
Displaying Message Filter Details	9-96
Configuring Filter Log Subscriptions	9-96
Changing Message Encoding	9-98
Sample Message Filters	9-100

Message Filter Examples	9-107
Open-Relay Prevention Filter	9-107
Policy Enforcement Filters	9-107
Routing and Domain Spoofing	9-111
Configuring Scan Behavior	9-115

CHAPTER 10**Mail Policies 10-1**

Overview of Mail Policies	10-1
How to Enforce Mail Policies on a Per-User Basis	10-2
Handling Incoming and Outgoing Messages Differently	10-3
Matching Users to a Mail Policy	10-3
First Match Wins	10-4
Examples of Policy Matching	10-4
Message Splintering	10-5
Managed Exceptions	10-6
Configuring Mail Policies	10-7
Configuring the Default Mail Policy for Incoming or Outgoing Messages	10-7
Creating a Mail Policy for a Group of Senders and Recipients	10-7
Finding Which Policies Apply to a Sender or Recipient	10-10

CHAPTER 11**Content Filters 11-1**

Overview of Content Filters	11-1
How Content Filters Work	11-1
How to Scan Message Content Using a Content Filter	11-2
Content Filter Conditions	11-2
Content Filter Actions	11-9
Action Variables	11-14
Filtering Messages Based on Content	11-16
Creating a Content Filter	11-16
Enabling Content Filters for All Recipients by Default	11-17
Applying the Content Filter to Messages for a Certain User Group	11-18
Notes on Configuring Content Filters in the GUI	11-18

CHAPTER 12**Anti-Virus 12-1**

Anti-Virus Scanning Overview	12-1
Evaluation Key	12-2
Scanning Messages with Multiple Anti-Virus Scanning Engines	12-2
Sophos Anti-Virus Filtering	12-2

Virus Detection Engine	12-3
Virus Scanning	12-3
Detection Methods	12-3
Virus Descriptions	12-4
Sophos Alerts	12-4
When a Virus is Found	12-4
McAfee Anti-Virus Filtering	12-5
Pattern-Matching Virus Signatures	12-5
Encrypted Polymorphic Virus Detection	12-5
Heuristics Analysis	12-5
When a Virus is Found	12-6
How to Configure the Appliance to Scan for Viruses	12-6
Enabling Virus Scanning and Configuring Global Settings	12-7
Configuring Virus Scanning Actions for Users	12-7
Configuring the Anti-Virus Policies for Different Groups of Senders and Recipients	12-13
Notes on Anti-Virus Configurations	12-14
Flow Diagram for Anti-Virus Actions	12-15
Sending an Email to the Appliance to Test Anti-Virus Scanning	12-16
Updating Virus Definitions	12-18
About Retrieving Anti-Virus Updates via HTTP	12-18
Configuring Update Server Settings	12-18
Monitoring and Manually Checking for Anti-Virus Updates	12-18
Verifying Anti-Virus Files Have Updated on the Appliance	12-20

CHAPTER 13

Anti-Spam 13-1

Overview of Anti-Spam Scanning	13-1
Anti-Spam Solutions	13-2
How to Configure the Appliance to Scan Messages for Spam	13-2
IronPort Anti-Spam Filtering	13-3
Evaluation Key	13-3
Cisco Anti-Spam: an Overview	13-4
Configuring IronPort Anti-Spam Scanning	13-5
Cisco Intelligent Multi-Scan Filtering	13-6
Configuring Cisco Intelligent Multi-Scan	13-7
Defining Anti-Spam Policies	13-7
Understanding Positive and Suspect Spam Thresholds	13-10
Configuration Examples: Actions for Positively Identified versus Suspected Spam	13-11
Unwanted Marketing Messages From Legitimate Sources	13-11

Using Custom Headers to Redirect URLs in Suspected Spam to the Cisco Web Security Proxy:
Configuration Example 13-11

Enabling Different Anti-Spam Scanning Engines in Different Mail Policies: Configuration
Example 13-12

Protecting Appliance-Generated Messages From the Spam Filter 13-14

Headers Added During Anti-Spam Scanning 13-14

Reporting Incorrectly Classified Messages to Cisco Systems 13-15

Determining Sender IP Address In Deployments with Incoming Relays 13-15

Example Environments with Incoming Relays 13-15

Configuring the Appliance to Work with Incoming Relays 13-17

How Incoming Relays Affect Functionality 13-22

Configuring Logs to Specify Which Headers Are Used 13-24

Monitoring Rules Updates 13-24

Testing Anti-Spam 13-25

Sending an Email to the Appliance to Test Cisco Anti-Spam 13-25

Ways Not to Test Anti-Spam Efficacy 13-26

CHAPTER 14

Outbreak Filters 14-1

Overview of Outbreak Filters 14-1

How Outbreak Filters Work 14-2

Delaying, Redirecting, and Modifying Messages 14-2

Threat Categories 14-2

Cisco Security Intelligence Operations 14-3

Context Adaptive Scanning Engine 14-4

Delaying Messages 14-4

Redirecting URLs 14-5

Modifying Messages 14-6

Types of Rules: Adaptive and Outbreak 14-6

Outbreaks 14-7

Threat Levels 14-7

How the Outbreak Filters Feature Works 14-8

Message Scoring 14-9

Dynamic Quarantine 14-10

Managing Outbreak Filters 14-11

Configuring Outbreak Filters Global Settings 14-12

Outbreak Filters Rules 14-15

The Outbreak Filters Feature and Mail Policies 14-15

The Outbreak Filters Feature and the Outbreak Quarantine 14-21

Monitoring Outbreak Filters 14-23

Outbreak Filters Report	14-23
Outbreak Filters Overview and Rules Listing	14-23
Outbreak Quarantine	14-23
Alerts, SNMP Traps, and Outbreak Filters	14-23
Troubleshooting The Outbreak Filters Feature	14-24
Reporting Incorrectly Classified Messages to Cisco	14-24
Multiple Attachments and Bypassed Filetypes	14-24
Message and Content Filters and the Email Pipeline	14-24

CHAPTER 15

URL Filtering 15-1

Overview of URL Filtering	15-1
Which URLs Are Evaluated	15-1
Setting Up URL Filtering	15-2
Requirements for URL Filtering	15-2
Enable URL Filtering	15-2
About the Connection to Cisco Web Security Services	15-3
URL Filtering in Cluster Configurations	15-4
Creating Whitelists for URL Filtering	15-4
Cisco Web Security Proxy End User Notification Page	15-5
Customizing the Appearance of End User Notification Page	15-6
Taking Action Based on the Reputation or Category of URLs in Messages	15-7
Using URL-Related Conditions (Rules) and Actions	15-7
Filtering by URL Reputation or URL Category: Conditions and Rules	15-8
Modifying URLs in Messages: Using URL Reputation and URL Category Actions in Filters	15-8
Redirected URLs: What Does the End User Experience?	15-10
Monitoring URL Filtering Results	15-10
Troubleshooting URL Filtering	15-10
Viewing Logs	15-11
Alert: SDS: Error Fetching Enrollment Certificate	15-11
Alert: SDS: Certificate Is Invalid	15-11
Unable to Connect to Cisco Web Security Services	15-11
Using the websecurityadvancedconfig Command	15-12
Message Tracking Search Does Not Find Messages with Specified Category	15-12
Malicious URLs and Marketing Messages Are Not Caught by Anti-Spam or Outbreak Filters	15-12
URLs in a Filtered Category Are Not Handled Correctly	15-13
End User Reaches Malicious Site via Rewritten URL	15-13
Manually Configuring a Certificate for Communication with Cisco Web Security Services	15-13
About URL Categories	15-13
URL Category Descriptions	15-14

Determining the Category of a URL	15-21
Reporting Uncategorized and Misclassified URLs	15-21
Future URL Category Set Changes	15-22

CHAPTER 16

File Reputation Filtering and File Analysis 16-1

Overview of File Reputation Filtering and File Analysis	16-1
File Threat Verdict Updates	16-1
File Processing Overview	16-2
Which Files Are Evaluated and Analyzed?	16-3
FIPS Compliance	16-4
Configuring File Reputation and Analysis Features	16-4
Requirements for Communication with File Reputation and Analysis Services	16-5
Enabling and Configuring File Reputation and Analysis Services	16-5
Configuring the Incoming Mail Policy for File Reputation Scanning and File Analysis	16-6
Quarantining Messages with Attachments Sent for Analysis	16-7
Using File Analysis Quarantine	16-8
Centralized File Analysis Quarantine	16-10
X-Headers for File Reputation and Analysis	16-10
Sending Notifications to End Users about Dropped Messages or Attachments	16-10
Advanced Malware Protection and Clusters	16-10
Ensuring That You Receive Alerts	16-11
Configuring Centralized Reporting for Advanced Malware Protection Features	16-11
File Reputation and File Analysis Reporting and Tracking	16-11
Identifying Files by SHA-256 Hash	16-11
File Reputation and File Analysis Report Pages	16-12
Viewing File Reputation Filtering Data in Other Reports	16-12
About Message Tracking and Advanced Malware Protection Features	16-12
Taking Action When File Threat Verdicts Change	16-13
Troubleshooting File Reputation and Analysis	16-13
Log Files	16-14
Using Trace	16-14
Several Alerts About Failure to Connect to File Reputation or File Analysis Servers	16-14
Many Files Have Verdict "Unscannable"	16-15

CHAPTER 17

Data Loss Prevention 17-1

Overview of Data Loss Prevention	17-1
Overview of the DLP Scanning Process	17-2
How Data Loss Prevention Works	17-2
DLP Deployment Options	17-3

System Requirements for Data Loss Prevention	17-4
RSA Email DLP	17-4
How to Set Up Data Loss Prevention for Deployments Using RSA Email DLP	17-4
Enabling Data Loss Prevention (RSA Email DLP)	17-5
DLP Policies for RSA Email DLP	17-6
DLP Policy Description	17-6
Predefined DLP Policy Templates	17-6
Setting Up RSA Email DLP Using a Wizard	17-7
Creating a DLP Policy Using a Predefined Template	17-8
Creating a Custom DLP Policy (Advanced)	17-9
About Defining Disallowed Content Using Content Matching Classifiers	17-10
Filtering Messages for DLP Policies	17-20
About Assessing Violation Severity	17-21
Arranging the Order of the Email DLP Policies for Violation Matching	17-21
Associating DLP Policies with Outgoing Mail Policies	17-22
Important Information About Editing or Deleting DLP Policies	17-23
RSA Enterprise Manager	17-23
How Enterprise Manager and the Email Security Appliance Work Together	17-24
Enterprise Manager Documentation	17-24
How to Set up Data Loss Prevention in Deployments with RSA Enterprise Manager	17-24
Migrating from RSA Email DLP to RSA Enterprise Manager	17-31
Checking for DLP Policy Updates from Enterprise Manager	17-32
RSA Enterprise Manager and Language Support	17-32
Using Enterprise Manager with Clustered Appliances	17-32
About Deleting and Disabling Policies in Enterprise Manager Deployments	17-33
Lost Connectivity Between the Email Security Appliance and Enterprise Manager	17-33
Switching from Enterprise Manager to RSA Email DLP	17-33
Message Actions	17-34
Defining Actions to Take for DLP Violations (Message Actions)	17-34
Viewing and Editing Message Actions	17-36
Drafting DLP Notifications	17-36
Showing or Hiding Sensitive DLP Data in Message Tracking	17-38
About Updating the DLP Engine and Content Matching Classifiers	17-39
Determining the Current Version of the RSA DLP Engine	17-39
Caveats for DLP Updates	17-40
Updating the DLP Engine and Content Matching Classifiers Manually	17-40
Enabling Automatic Updates (Not Recommended)	17-40
DLP Updates on Centralized (Clustered) Appliances	17-41
Rolling Back DLP Updates	17-41

Working with DLP Incident Messages and Data	17-41
Troubleshooting Data Loss Prevention	17-42
Enterprise Manager Disconnects the Email Security Appliance	17-42
RSA Email DLP Fails to Detect Violations in Email Attachments	17-43

CHAPTER 18**Cisco Email Encryption 18-1**

Overview of Cisco Email Encryption	18-1
How to Encrypt Messages with a Local Key Server	18-2
Encryption Workflow	18-2
Encrypting Messages using the Email Security Appliance	18-4
Enabling Message Encryption on the Email Security Appliance	18-4
Configuring How a Key Service Handles Encrypted Messages	18-4
Configuring the Default Locale of the Envelope	18-7
Updating to the Latest Version of the PXE Engine	18-8
Determining Which Messages to Encrypt	18-8
Using a TLS Connection as an Alternative to Encryption	18-9
Encrypting and Immediately Delivering Messages using a Content Filter	18-9
Encrypting a Message upon Delivery using a Content Filter	18-10
Inserting Encryption Headers into Messages	18-11
Encryption Headers	18-12
Encryption Headers Examples	18-14

CHAPTER 19**S/MIME Security Services 19-1**

Overview of S/MIME Security Services	19-1
S/MIME Security Services in Email Security Appliance	19-1
Understanding How S/MIME Security Services Works	19-2
Signing, Encrypting, or Signing and Encrypting Outgoing Messages using S/MIME	19-4
S/MIME Signing and Encryption Workflow in Email Security Appliance	19-4
How to Sign, Encrypt, or Sign and Encrypt Outgoing Messages using S/MIME	19-6
Setting Up Certificates for S/MIME Signing	19-6
Setting Up Public Keys for S/MIME Encryption	19-9
Managing S/MIME Sending Profiles	19-10
Determining Which Messages to Sign, Encrypt, or Sign and Encrypt	19-13
Signing, Encrypting, or Signing and Encrypting and Immediately Delivering Messages using a Content Filter	19-13
Signing, Encrypting, or Signing and Encrypting a Message upon Delivery using a Content Filter	19-14
Verifying, Decrypting, or Decrypting and Verifying Incoming Messages using S/MIME	19-14
S/MIME Verification and Decryption Workflow in Email Security Appliance	19-15

How to Verify, Decrypt, or Decrypt and Verify Incoming Messages Using S/MIME	19-16
Setting Up Certificates for Decrypting Messages	19-16
Setting Up Public Keys for Verifying Signed Messages	19-17
Enabling S/MIME Decryption and Verification	19-19
Configuring an Action for S/MIME Decrypted or Verified Message	19-20
S/MIME Certificate Requirements	19-20
Certificate Requirements for Signing	19-20
Certificate Requirements for Encryption	19-21
Exporting Public Keys	19-23

CHAPTER 20

Email Authentication 20-1

Email Authentication Overview	20-1
DomainKeys and DKIM Authentication	20-1
DomainKeys and DKIM Authentication Workflow	20-2
DomainKeys and DKIM Signing in AsyncOS	20-2
Configuring DomainKeys and DKIM Signing	20-3
Signing Keys	20-3
Public Keys	20-4
Domain Profiles	20-5
Enabling Signing for Outgoing Mail	20-6
Enabling Signing for Bounce and Delay Messages	20-6
Configuring DomainKeys/DKIM Signing (GUI)	20-7
Domain Keys and Logging	20-16
How to Verify Incoming Messages Using DKIM	20-16
DKIM Verification Checks Performed by AsyncOS	20-17
Managing DKIM Verification Profiles	20-17
Configuring DKIM Verification on the Mail Flow Policy	20-20
Configuring an Action for DKIM Verified Mail	20-21
Overview of SPF and SIDF Verification	20-22
How to Verify Incoming Messages Using SPF/SIDF	20-23
Enabling SPF and SIDF	20-24
Determining the Action to Take for SPF/SIDF Verified Mail	20-31
Verification Results	20-31
Using the spf-status Filter Rule in the CLI	20-32
spf-status Content Filter Rule in the GUI	20-33
Using the spf-passed Filter Rule	20-33
Testing the SPF/SIDF Results	20-34
Basic Granularity Test of SPF/SIDF Results	20-34

Greater Granularity Test of SPF/SIDF Results	20-34
DMARC Verification	20-35
DMARC Verification Workflow in AsyncOS for Email	20-36
How to Verify Incoming Messages Using DMARC	20-36

CHAPTER 21

Text Resources 21-1

Overview of Text Resources	21-1
Content Dictionaries	21-1
Text Resources	21-2
Message Disclaimer Stamping	21-2
Content Dictionaries	21-2
Dictionary Content	21-2
Importing and Exporting Dictionaries as Text Files	21-3
Adding Dictionaries	21-4
Deleting Dictionaries	21-5
Importing Dictionaries	21-5
Exporting Dictionaries	21-5
Using and Testing the Content Dictionaries Filter Rules	21-6
Dictionary Match Filter Rule	21-6
Understanding Text Resources	21-8
Importing and Exporting Text Resources as Text Files	21-8
Overview of Text Resource Management	21-9
Adding Text Resources	21-9
Deleting Text Resources	21-10
Importing Text Resources	21-10
Exporting Text Resources	21-10
Overview of HTML-Based Text Resources	21-11
Using Text Resources	21-12
Disclaimer Template	21-12
Disclaimer Stamping and Multiple Encodings	21-16
Notification Templates	21-19
Anti-Virus Notification Templates	21-20
Bounce and Encryption Failure Notification Templates	21-23
Encryption Notification Templates	21-24

CHAPTER 22

Validating Recipients Using an SMTP Server 22-1

Overview of SMTP Call-Ahead Recipient Validation	22-1
SMTP Call-Ahead Recipient Validation Workflow	22-1

How to Validate Recipients Using an External SMTP Server	22-3
Configuring the Call-Ahead Server Profile	22-3
Enabling a Listener to Validate Incoming Mail Via the SMTP Server	22-6
Configuring LDAP Routing Query Settings	22-6
SMTP Call-Ahead Query Routing	22-7
Bypassing SMTP Call-Ahead Validation for Certain Users or Groups	22-8

CHAPTER 23**Encrypting Communication with Other MTAs 23-1**

Overview of Encrypting Communication with Other MTAs	23-1
How to Encrypt SMTP Conversations using TLS	23-2
Obtaining Certificates	23-2
Intermediate Certificates	23-3
Certificates and Centralized Management	23-3
Creating a Self-Signed Certificate using the GUI	23-3
Importing a Certificate Using the GUI	23-5
Creating a Self-Signed Certificate or Importing a Certificate using the CLI	23-6
Exporting a Certificate Using the GUI	23-6
Enabling TLS on a Listener's HAT	23-6
Assigning a Certificate to a Public or Private Listener for TLS Connections Using the GUI	23-7
Assigning a Certificate to a Public or Private Listener for TLS Connections Using the CLI	23-8
Logging	23-8
GUI Example: Changing the TLS Setting for Listener's HAT	23-8
CLI Example: Changing the TLS Setting for Listener's HAT	23-9
Enabling TLS and Certificate Verification on Delivery	23-10
Sending Alerts When a Required TLS Connection Fails	23-11
Logging	23-12
CLI Example	23-12
Managing Lists of Certificate Authorities	23-16
Viewing the Pre-Installed list of Certificate Authorities	23-17
Disabling the System Certificate Authority List	23-17
Importing a Custom Certificate Authority List	23-17
Exporting a Certificate Authorities List	23-18
Enabling a Certificate for HTTPS	23-18

CHAPTER 24**Configuring Routing and Delivery Features 24-1**

Routing Email for Local Domains	24-1
SMTP Routes Overview	24-2
Default SMTP Route	24-2

Defining an SMTP Route	24-3
SMTP Routes Limits	24-3
SMTP Routes and DNS	24-3
SMTP Routes and Alerts	24-4
SMTP Routes, Mail Delivery, and Message Splintering	24-4
SMTP Routes and Outbound SMTP Authentication	24-4
Managing SMTP Routes to Send Outbound Email Using the GUI	24-4
Rewriting Addresses	24-6
Creating Alias Tables	24-7
Configuring an Alias Table from the Command Line	24-7
Exporting and Importing an Alias Table	24-8
Deleting Entries from the Alias Table	24-9
Configuring Masquerading	24-16
Masquerading and altsrghost	24-17
The Domain Map Feature	24-28
Importing and Exporting a Domain Map Table	24-34
Directing Bounced Email	24-35
Handling Undeliverable Email	24-36
Creating a New Bounce Profile	24-40
Applying Bounce Profiles to Listeners	24-41
Controlling Email Delivery Using Destination Controls	24-42
Determining Which Interface is Used for Mail Delivery	24-43
Default Delivery Limits	24-44
Working with Destination Controls	24-44
Bounce Verification	24-51
Overview: Tagging and Bounce Verification	24-52
Accepting Legitimate Untagged Bounced Messages	24-53
Preventing a Bounced Message Storm Using Bounce Verification	24-54
Set Email Delivery Parameters	24-56
Default Delivery IP Interface	24-56
Possible Delivery Feature	24-56
Default Maximum Concurrency	24-57
deliveryconfig Example	24-57
Configuring Mail Gateways for all Hosted Domains Using Virtual Gateway™ Technology	24-59
Overview	24-60
Setting Up Virtual Gateway Addresses	24-60
Monitoring the Virtual Gateway Addresses	24-68
Managing Delivery Connections per Virtual Gateway Address	24-68
Using Global Unsubscribe	24-69

Adding a Global Unsubscribe Address Using The CLI	24-70
Exporting and Importing a Global Unsubscribe File	24-72
Review: Email Pipeline	24-74

CHAPTER 25

LDAP Queries 25-1

Overview of LDAP Queries	25-1
Understanding LDAP Queries	25-2
Understanding How LDAP Works with AsyncOS	25-3
Configuring the Cisco IronPort Appliance to Work with an LDAP Server	25-4
Creating LDAP Server Profiles to Store Information About the LDAP Server	25-5
Testing LDAP Servers	25-6
Enabling LDAP Queries to Run on a Particular Listener	25-7
Enhanced Support for Microsoft Exchange 5.5	25-9
Working with LDAP Queries	25-12
Types of LDAP Queries	25-12
Base Distinguishing Name (DN)	25-13
LDAP Query Syntax	25-13
Secure LDAP (SSL)	25-14
Routing Queries	25-14
Allowing Clients to Bind to the LDAP Server Anonymously	25-14
Testing LDAP Queries	25-17
Troubleshooting Connections to LDAP Servers	25-18
Using Acceptance Queries For Recipient Validation	25-19
Sample Acceptance Queries	25-19
Configuring Acceptance Queries for Lotus Notes	25-20
Using Routing Queries to Send Mail to Multiple Target Addresses	25-20
Sample Routing Queries	25-21
Using Masquerading Queries to Rewrite the Envelope Sender	25-21
Sample Masquerading Queries	25-22
Masquerading "Friendly Names"	25-22
Using Group LDAP Queries to Determine if a Recipient is a Group Member	25-23
Sample Group Queries	25-23
Configuring a Group Query	25-23
Example: Using a Group Query to Skip Spam and Virus Checking	25-25
Using Domain-based Queries to Route to a Particular Domain	25-26
Creating a Domain-Based Query	25-27
Using Chain Queries to Perform a Series of LDAP Queries	25-28
Creating a Chain Query	25-28

Using LDAP For Directory Harvest Attack Prevention	25-29
Directory Harvest Attack Prevention within the SMTP Conversation	25-29
Directory Harvest Attack Prevention within the Work Queue	25-31
Configuring AsyncOS for SMTP Authentication	25-32
Configuring SMTP Authentication	25-33
Configuring an SMTP Authentication Query	25-34
SMTP Authentication via Second SMTP Server (SMTP Auth with Forwarding)	25-35
SMTP Authentication with LDAP	25-36
Authenticating SMTP Sessions Using Client Certificates	25-39
Outgoing SMTP Authentication	25-39
Logging and SMTP Authentication	25-40
Configuring External LDAP Authentication for Users	25-40
User Accounts Query	25-41
Group Membership Queries	25-41
Authenticating End-Users of the Spam Quarantine	25-43
Sample Active Directory End-User Authentication Settings	25-43
Sample OpenLDAP End-User Authentication Settings	25-44
Spam Quarantine Alias Consolidation Queries	25-44
Sample Active Directory Alias Consolidation Settings	25-45
Sample OpenLDAP Alias Consolidation Settings	25-45
Identifying a Sender's User Distinguished Name for RSA Enterprise Manager	25-45
Sample User Distinguished Name Settings	25-46
Configuring AsyncOS To Work With Multiple LDAP Servers	25-46
Testing Servers and Queries	25-47
Failover	25-47
Load Balancing	25-48

CHAPTER 26

Authenticating SMTP Sessions Using Client Certificates	26-49
Overview of Certificates and SMTP Authentication	26-49
How to Authenticate a User with a Client Certificate	26-50
How to Authenticate a User with an SMTP Authentication LDAP Query	26-50
How to Authenticate a User with an LDAP SMTP Authentication Query if the Client Certificate is Invalid	26-50
Checking the Validity of a Client Certificate	26-51
Authenticating a User Using an LDAP Directory	26-52
Authenticating an SMTP Connection Over TLS Using a Client Certificate	26-52
Establishing a TLS Connection from the Appliance	26-53
Updating a List of Revoked Certificates	26-54

CHAPTER 27**FIPS Management 27-1**

- FIPS Management Overview 27-1
- Configuration Changes in FIPS Mode 27-1
- Switching the Appliance to FIPS Mode 27-2
- Encrypting Sensitive Data in FIPS Mode 27-3
- Checking FIPS Mode Compliance 27-4
- Managing Certificates and Keys 27-4
- Managing Keys for DKIM Signing and Verification 27-5
 - DKIM Signing 27-5
 - DKIM Verification 27-6

CHAPTER 28**Using Email Security Monitor 28-1**

- Email Security Monitor Overview 28-1
 - Email Security Monitor and Centralized Management 28-2
- Email Security Monitor Pages 28-2
 - Searching and Email Security Monitor 28-4
 - Viewing Details of Messages Included in Reports 28-4
 - My Reports Page 28-5
 - Overview Page 28-6
 - Incoming Mail Page 28-9
 - Outgoing Destinations 28-14
 - Outgoing Senders 28-15
 - Delivery Status Page 28-15
 - Internal Users Page 28-16
 - DLP Incidents Page 28-17
 - Content Filters Page 28-18
 - DMARC Verification Page 28-19
 - Outbreak Filters Page 28-19
 - Virus Types Page 28-21
 - URL Filtering Page 28-21
 - File Reputation and File Analysis Reports 28-22
 - TLS Connections Page 28-22
 - Inbound SMTP Authentication Page 28-23
 - Rate Limits Page 28-23
 - System Capacity Page 28-24
 - System Status Page 28-30
 - High Volume Mail Page 28-32
 - Message Filters Page 28-32

Retrieving CSV Data	28-33
Reporting Overview	28-35
Scheduled Report Types	28-35
Setting the Return Address for Reports	28-36
Managing Reports	28-36
Scheduled Reports	28-36
Archived Reports	28-38
Troubleshooting Email Reports	28-39

CHAPTER 29

Tracking Messages	29-1
Message Tracking Overview	29-1
Enabling Message Tracking	29-1
Searching for Messages	29-2
Working with Message Tracking Search Results	29-4
Message Details	29-5
Checking Message Tracking Data Availability	29-6
About Message Tracking and Upgrades	29-6
Troubleshooting Message Tracking	29-7
Attachments Do Not Appear in Search Results	29-7
Expected Messages Are Missing from Search Results	29-7

CHAPTER 30

Policy, Virus, and Outbreak Quarantines	30-1
Overview of Policy, Virus, and Outbreak Quarantines	30-1
Quarantine Types	30-2
Managing Policy, Virus, and Outbreak Quarantines	30-3
Disk Space Allocation for Policy, Virus, and Outbreak Quarantines	30-3
Retention Time for Messages in Quarantines	30-3
Default Actions for Automatically Processed Quarantined Messages	30-4
Checking the Settings of System-Created Quarantines	30-5
Configuring Policy, Virus, and Outbreak Quarantines	30-5
About Editing Policy, Virus, and Outbreak Quarantine Settings	30-6
Determining the Filters and Message Actions to Which a Policy Quarantine Is Assigned	30-7
About Deleting Policy Quarantines	30-7
Monitoring Quarantine Status, Capacity, and Activity	30-8
Policy Quarantine Performance	30-8
Alerts About Quarantine Disk-Space Usage	30-9
Policy Quarantines and Logging	30-9
About Distributing Message Processing Tasks to Other Users	30-9

About Policy, Virus, and Outbreak Quarantines in Cluster Configurations	30-10
About Centralized Policy, Virus, and Outbreak Quarantines	30-10
Working with Messages in Policy, Virus, or Outbreak Quarantines	30-10
Viewing Messages in Quarantines	30-11
Finding Messages in Policy, Virus, and Outbreak Quarantines	30-11
Manually Processing Messages in a Quarantine	30-12
Messages in Multiple Quarantines	30-13
Message Details and Viewing Message Content	30-14
About Rescanning of Quarantined Messages	30-17
The Outbreak Quarantine	30-17

CHAPTER 31

Spam Quarantine 31-1

Overview of the Spam Quarantine	31-1
Local Versus External Spam Quarantine	31-1
Setting Up the Local Spam Quarantine	31-2
Enabling and Configuring the Spam Quarantine	31-3
Configuring the IP Interface for Browser Access to the Spam Quarantine	31-4
Configuring Administrative User Access to the Spam Quarantine	31-4
Configuring a Mail Policy to Quarantine Spam	31-5
Limiting Which Recipients Have Mail Quarantined	31-5
Ensuring That Message Text Displays Correctly	31-6
Spam Quarantine Language	31-6
Using Safelists and Blocklists to Control Email Delivery Based on Sender	31-7
Message Processing of Safelists and Blocklists	31-7
Enabling Safelists and Blocklists	31-8
External Spam Quarantine and Safelist/Blocklists	31-8
Adding Senders and Domains to Safelists and Blocklists (Administrators)	31-9
About End-User Access to Safelists and Blocklists	31-11
Synchronizing Safelists/Blocklists on Multiple Email Security Appliances (Deployments Without a Security Management Appliance)	31-12
Backing Up and Restoring the Safelist/Blocklist	31-13
Troubleshooting Safelists and Blocklists	31-13
Configuring Spam Management Features for End Users	31-14
Authentication Options for End Users Accessing Spam Management Features	31-15
Setting Up End-User Access to the Spam Quarantine via Web Browser	31-16
Notifying End Users About Quarantined Messages	31-19
Managing Messages in the Spam Quarantine	31-22
Accessing the Spam Quarantine (Administrative Users)	31-22
Searching for Messages in the Spam Quarantine	31-22

Viewing Messages in the Spam Quarantine	31-23
Delivering Messages in the Spam Quarantine	31-23
Deleting Messages from the Spam Quarantine	31-24
Disk Space for the Spam Quarantine	31-24
About Disabling the Spam Quarantine	31-24
Troubleshooting Spam Quarantine Features	31-24

CHAPTER 32**Distributing Administrative Tasks 32-1**

Working with User Accounts	32-1
User Roles	32-2
Managing Users	32-3
Managing Custom User Roles for Delegated Administration	32-7
Account Privileges Page	32-8
Assigning Access Privileges	32-9
Defining a Custom User Role	32-13
Defining a Custom User Role When Adding a User Account	32-14
Updating Responsibilities for a Custom User Role	32-14
Editing a Custom User Role	32-15
Duplicating a Custom User Role	32-15
Deleting a Custom User Role	32-16
Passwords	32-16
Changing Your Password	32-16
Locking and Unlocking a User Account	32-17
Configuring Restrictive User Account and Password Settings	32-17
External Authentication	32-20
Configuring Access to the Email Security Appliance	32-24
Configuring IP-Based Network Access	32-24
Configuring Session Timeouts	32-26
Displaying Messages to Administrative Users	32-27
Managing Secure Shell (SSH) Keys	32-28
Example: Install a New Public Key	32-28
Example: Edit SSH Server Configuration	32-29
Remote SSH Command Execution	32-30
Viewing Active Administrator Sessions	32-30

CHAPTER 33**System Administration 33-1**

Management of the Appliance	33-1
Shutting Down or Rebooting the Appliance	33-2

Suspending Email Receiving and Delivery	33-2
Resuming Suspended Email Receiving and Delivery	33-3
Taking an Appliance Offline Using the CLI	33-3
Resetting to Factory Defaults	33-4
Displaying the Version Information for AsyncOS	33-5
Feature Keys	33-5
Adding and Managing Feature Keys	33-5
Automating Feature Key Download and Activation	33-6
Expired Feature Keys	33-7
Cisco Email Security Virtual Appliance License	33-7
Virtual Appliance License Expiration	33-7
Managing the Configuration File	33-7
Managing Configuration Files Using the GUI	33-8
CLI Commands for Configuration Files	33-12
Managing Disk Space	33-16
(Virtual Appliances Only) Increasing Available Disk Space	33-16
Allocating Disk Space Quotas	33-16
Managing Disk Space for the Miscellaneous Quota	33-17
Ensuring That You Receive Alerts About Disk Space	33-17
Disk Space and Centralized Management	33-17
Service Updates	33-17
Setting Up to Obtain Upgrades and Updates	33-18
Options for Distributing Upgrades and Updates	33-18
Configuring Your Network to Download Upgrades and Updates from the Cisco Servers	33-19
Configuring the Appliance for Upgrades and Updates in Strict Firewall Environments	33-19
Upgrading and Updating from a Local Server	33-19
Hardware and Software Requirements for Upgrading and Updating from a Local Server	33-20
Hosting an Upgrade Image on a Local Server	33-21
UpdatesThrough a Proxy Server	33-21
Configuring Server Settings for Downloading Upgrades and Updates	33-21
Configuring Automatic Updates	33-23
Configuring the Appliance to Verify the Validity of Updater Server Certificate	33-23
Configuring the Appliance to Trust Proxy Server Communication	33-25
Upgrading AsyncOS	33-25
About Upgrading Clustered Systems	33-26
About Batch Commands for Upgrade Procedures	33-26
Notifications of Available Upgrades	33-26
Preparing to Upgrade AsyncOS	33-26
Downloading and Installing the Upgrade	33-27

Enabling Remote Power Management	33-29
Reverting to a Previous Version of AsyncOS	33-30
Reversion Impact	33-30
Reverting AsyncOS	33-31
Configuring the Return Address for Appliance Generated Messages	33-33
Alerts	33-34
AutoSupport	33-34
Alert Delivery	33-34
Adding Alert Recipients	33-36
Configuring Alert Settings	33-37
Viewing Recent Alerts	33-38
Alert Descriptions	33-38
Changing Network Settings	33-53
Changing the System Hostname	33-53
Configuring Domain Name System (DNS) Settings	33-54
Configuring TCP/IP Traffic Routes	33-57
Configuring the Default Gateway	33-57
Configuring SSL Settings	33-58
Disabling SSLv3 for Enhanced Security	33-58
System Time	33-59
Selecting a Time Zone	33-59
Editing Time Settings	33-60
Customizing Your View	33-60
Using Favorite Pages	33-60
Setting User Preferences	33-61
Overriding Internet Explorer Compatibility Mode	33-62

CHAPTER 34

Managing and Monitoring Using the CLI	34-1
Overview of Managing and Monitoring Using the CLI	34-1
Reading the Available Components of Monitoring	34-2
Reading the Event Counters	34-2
Reading the System Gauges	34-4
Reading the Rates of Delivered and Bounced Messages	34-6
Monitoring Using the CLI	34-6
Monitoring the Email Status	34-7
Monitoring Detailed Email Status	34-8
Monitoring the Status of a Mail Host	34-11
Determining the Make-up of the Email Queue	34-15

Displaying Real-time Activity	34-16
Monitoring Inbound Email Connections	34-19
Checking the DNS Status	34-20
Resetting Email Monitoring Counters	34-21
Identifying Active TCP/IP Services	34-22
Managing the Email Queue	34-22
Deleting Recipients in Queue	34-22
Bouncing Recipients in Queue	34-24
Redirecting Messages in Queue	34-26
Showing Messages Based on Recipient in Queue	34-27
Suspending Email Delivery	34-28
Resuming Email Delivery	34-29
Suspending Receiving Email	34-29
Resuming Receiving Email	34-30
Resuming Delivery and Receiving of Email	34-31
Scheduling Email for Immediate Delivery	34-31
Pausing the Work Queue	34-32
Locating and Archiving Older Messages	34-34
Tracking Messages Within the System	34-35
SNMP Monitoring	34-36
MIB Files	34-37
Hardware Objects	34-37
SNMP Traps	34-39

CHAPTER 35

SenderBase Network Participation	35-1
Overview of SenderBase Network Participation	35-1
Sharing Statistics with SenderBase	35-1
Frequently Asked Questions	35-2

CHAPTER 36

Other Tasks in the GUI	36-7
The Graphical User Interface (GUI)	36-7
Enabling the GUI on an Interface	36-7
System Information in the GUI	36-11
Gathering XML status from the GUI	36-11

CHAPTER 37

Advanced Network Configuration	37-1
Media Settings on Ethernet Interfaces	37-1
Using etherconfig to Edit Media Settings on Ethernet Interfaces	37-1

Network Interface Card Pairing/Teaming	37-3
NIC Pair Naming	37-4
Virtual Local Area Networks (VLANs)	37-6
VLANs and Physical Ports	37-7
Managing VLANs	37-8
Direct Server Return	37-13
Enabling Direct Server Return	37-13
Ethernet Interface's Maximum Transmission Unit	37-18

CHAPTER 38
Logging 38-1

Overview	38-1
Understanding Log Files and Log Subscriptions	38-1
Log Types	38-1
Log Retrieval Methods	38-6
Log Types	38-7
Timestamps in Log Files	38-8
Using Text Mail Logs	38-8
Using Delivery Logs	38-15
Using Bounce Logs	38-17
Using Status Logs	38-19
Using Domain Debug Logs	38-22
Using Injection Debug Logs	38-23
Using System Logs	38-24
Using CLI Audit Logs	38-25
Using FTP Server Logs	38-26
Using HTTP Logs	38-27
Using NTP Logs	38-28
Using Scanning Logs	38-28
Using Anti-Spam Logs	38-29
Using Anti-Virus Logs	38-29
Using Spam Quarantine Logs	38-30
Using Spam Quarantine GUI Logs	38-30
Using LDAP Debug Logs	38-31
Using Safelist/Blocklist Logs	38-32
Using Reporting Logs	38-33
Using Reporting Query Logs	38-34
Using Updater Logs	38-35
Understanding Tracking Logs	38-36
Using Authentication Logs	38-37

Using Configuration History Logs	38-37
Log Subscriptions	38-38
Configuring Log Subscriptions	38-39
Creating a Log Subscription in the GUI	38-40
Configuring Global Settings for Logging	38-40
Rolling Over Log Subscriptions	38-43
Viewing Recent Log Entries in the GUI	38-45
Viewing Recent Log Entries in the CLI (tail Command)	38-45
Configuring Host Keys	38-47

CHAPTER 39

Centralized Management Using Clusters	39-1
Overview of Centralized Management Using Clusters	39-1
Cluster Requirements	39-2
Cluster Organization	39-2
Initial Configuration Settings	39-3
Creating and Joining a Cluster	39-4
The clusterconfig Command	39-4
Adding Groups	39-10
Managing Clusters	39-10
Administering a Cluster from the CLI	39-10
Copying and Moving Settings	39-11
Experimenting with New Configurations	39-11
Leaving a Cluster Permanently (Removal)	39-12
Upgrading Machines in a Cluster	39-12
CLI Command Support	39-13
All Commands Are Cluster-aware	39-13
Restricted Commands	39-14
Administering a Cluster from the GUI	39-15
Cluster Communication	39-18
DNS and Hostname Resolution	39-18
Cluster Communication Security	39-19
Cluster Consistency	39-19
Disconnect/Reconnect	39-20
Interdependent Settings	39-21
Loading a Configuration in Clustered Appliances	39-22
Best Practices and Frequently Asked Questions	39-24
Best Practices	39-24
Setup and Configuration Questions	39-27

General Questions	39-28
Network Questions	39-28
Planning and Configuration	39-29

CHAPTER 40**Testing and Troubleshooting 40-1**

Debugging Mail Flow Using Test Messages: Trace	40-1
Using the Listener to Test the Appliance	40-12
Troubleshooting the Network	40-16
Testing the Network Connectivity of the Appliance	40-16
Troubleshooting the Listener	40-22
Troubleshooting Email Delivery From the Appliance	40-23
Troubleshooting Performance	40-26
Responding to Alerts	40-27
Alert: Battery Relearn Timed Out (RAID Event) on C380 or C680 Hardware	40-27
Troubleshooting Alerts That Miscellaneous Disk Usage is Approaching the Quota	40-27
Remotely Resetting Appliance Power	40-27
Working with Technical Support	40-28
Technical Support for Virtual Appliances	40-28
Opening or Updating a Support Case From the Appliance	40-28
Enabling Remote Access for Cisco Technical Support Personnel	40-29
Running a Packet Capture	40-31

CHAPTER 41**Optimizing the Appliance for Outbound Mail Delivery Using D-Mode 41-1**

Feature Summary: D-Mode for Optimized Outbound Delivery	41-1
Features Unique to D-Mode-Enabled Appliances	41-1
Standard Features Disabled in D-Mode-Enabled Appliances	41-2
Standard Features Applicable to D-Mode-Enabled Appliances	41-2
Setting Up the Appliance for Optimized Outbound Mail Delivery	41-3
Configuring Resource-Conserving Bounce Settings	41-3
Sending Bulk Mail Using IronPort Mail Merge (IPMM)	41-4
Overview of IronPort Mail Merge	41-4
Benefits of the Mail Merge Function	41-5
Using Mail Merge	41-5
Command Descriptions	41-8
Notes on Defining Variables	41-8
Example IPMM Conversation	41-9

CHAPTER 42**Centralizing Services on a Cisco Content Security Management Appliance 42-1**

- Overview of Cisco Content Security Management Appliance Services 42-1
- Network Planning 42-2
- Working with an External Spam Quarantine 42-2
 - Mail Flow and the External Spam Quarantine 42-2
 - Migrating from a Local Spam Quarantine to an External Quarantine 42-3
 - Enabling an External Spam Quarantine and External Safelist/Blocklist 42-3
 - Disabling the Local Spam Quarantine to Activate the External Quarantine 42-4
 - Troubleshooting an External Spam Quarantine 42-5
- About Centralizing Policy, Virus, and Outbreak Quarantines 42-5
 - Centralized Policy, Virus, and Outbreak Quarantines 42-5
 - About Migration of Policy, Virus, and Outbreak Quarantines 42-6
 - Centralizing Policy, Virus, and Outbreak Quarantines 42-7
 - About Disabling Centralized Policy, Virus, and Outbreak Quarantines 42-8
 - Troubleshooting Centralized Policy, Virus, and Outbreak Quarantines 42-9
- Configuring Centralized Reporting 42-10
- Configuring Centralized Message Tracking 42-11
- Using Centralized Services 42-11

APPENDIX A**FTP, SSH, SCP, and Telnet Access A-1**

- IP Interfaces A-1
 - How AsyncOS Selects Default IP Interface A-2
- Configuring FTP Access to the Email Security Appliance A-2
- Secure Copy (scp) Access A-5
- Accessing the Email Security appliance via a Serial Connection A-5

APPENDIX B**Assigning Network and IP Addresses B-1**

- Ethernet Interfaces B-1
- Selecting IP Addresses and Netmasks B-1
 - Sample Interface Configurations B-2
 - IP Addresses, Interfaces, and Routing B-3
 - Summary B-3
- Strategies for Connecting Your Cisco Appliance B-3

CHAPTER C**Example of Mail Policies and Content Filters C-1**

- Overview of Incoming Mail Policies C-1
- Accessing Mail Policies C-1

Configuring the Default Anti-Spam Policies for Incoming Messages	C-3
Creating a Mail Policy for a Group of Sender and Recipients	C-4
Creating Mail Policies for Different Groups of Senders and Recipients	C-7
Finding Senders or Recipients in Mail Policies	C-10
Filtering Messages Based on Content	C-12
Applying Individual Content Filters to Different Groups of Recipients	C-15
Notes on Configuring Content Filters in the GUI	C-17

APPENDIX D

Firewall Information D-1

APPENDIX E

End User License Agreement E-1

Cisco Systems End User License Agreement	E-1
Supplemental End User License Agreement for Cisco Systems Content Security Software	E-8

GLOSSARY

INDEX



Getting Started with the Cisco Email Security Appliance

- [What's New in This Release, page 1-1](#)
- [Where to Find More Information, page 1-7](#)
- [Cisco Email Security Appliance Overview, page 1-10](#)

What's New in This Release

- [What's New in Cisco AsyncOS 9.1 for Email, page 1-2](#)
- [What's New in Cisco AsyncOS 9.0 for Email, page 1-3](#)

What's New in Cisco AsyncOS 9.1 for Email

Feature	Description
New Features	
File Analysis quarantine improvements	Messages can now be automatically released or deleted from the centralized File Analysis quarantine on the Content Security Management Appliance based on the analysis verdict.
Updater enhancements	Email Security appliance includes the following updater enhancements: - Email security appliance can check the validity of the Cisco updater server certificate every time the appliance communicates with the updater server. See Configuring the Appliance to Verify the Validity of Updater Server Certificate, page 33-23 - If you are using a non-transparent proxy server, you can add the CA certificate used to sign the proxy certificate to the appliance. By doing so, the appliance trusts the proxy server communication. See Configuring the Appliance to Trust Proxy Server Communication, page 33-25.
Disabling SSLv3	For enhanced security, you can disable SSLv3 for the following services: - Updater - URL Filtering - End User Quarantine - LDAP See Disabling SSLv3 for Enhanced Security, page 33-58.
Changed Behavior	
Accept or reject messages based on the size of the subject	While configuring the global settings for listeners, you can now specify whether to accept or reject messages based on the size of the subject. If you specify this parameter, messages having subject size within the specified limit will be accepted and any other messages will be rejected. For instructions, see Configuring Global Settings for Listeners , page 5-6.

What's New in Cisco AsyncOS 9.0 for Email

Feature	Description
New Features	
Virtual Appliance Enhancements	The following are now supported: • Access to more than 2 TB of disk space, for virtual appliances running ESXi 5.5 and VMFS 5 • ESXi 5.5 hypervisor • Thin provisioning In addition, feature keys that expire independently of the virtual appliance license are now available, allowing evaluation feature keys to be deployed on virtual appliances.
Release and Support Notifications	You can now receive software release and critical support notifications from Cisco Support (in the form of alerts). See Adding Alert Recipients, page 33-36.
S/MIME Security Services	AsyncOS for Email now allows organizations to communicate securely using S/MIME without requiring that all end-users possess their own certificates. Organizations can handle message signing, encryption, verification, and decryption at the gateway level using certificates that identify the organization rather than the individual. AsyncOS provides the following S/MIME security services: • Sign, encrypt, or sign and encrypt messages using S/MIME. See Signing, Encrypting, or Signing and Encrypting Outgoing Messages using S/MIME, page 19-4. • Verify, decrypt, or decrypt and verify messages using S/MIME. See Verifying, Decrypting, or Decrypting and Verifying Incoming Messages using S/MIME, page 19-14.
Cisco AsyncOS API for Email	The Cisco AsyncOS API for Email (or AsyncOS API) is a Representational StateTransfer (REST)-based set of operations that provide secure and authenticated access to the Email Security appliance reports and report counters. You can retrieve the Email Security appliance reporting data using this API. See <i>Cisco AsyncOS API for Email - Getting Started Guide</i>.
Advanced Malware Protection Quarantine	AsyncOS for Email now includes an Advanced Malware Protection-specific quarantine. You can configure the appliance to quarantine messages with attachments sent for analysis. See Quarantining Messages with Attachments Sent for Analysis, page 16-7.
Enhancements	

Feature	Description
Advanced Malware Protection Improvements	- You can now use the Advanced Malware Protection feature to detect malware in archived or compressed email attachments. For more information, see Archive or Compressed File Processing, page 16-4. For a list of supported archive and compressed formats, see <i>File Criteria for Advanced Malware Protection Services for Cisco Content Security Products</i>, available from http://www.cisco.com/c/en/us/support/security/email-security-appliance/products-user-guide-list.html. - When you configure the file analysis feature, you choose which file types are sent for analysis. See Enabling and Configuring File Reputation and Analysis Services, page 16-5. - New types are added dynamically; you will receive an alert when the list of uploadable file types changes, and can select added file types to upload. See Ensuring That You Receive Alerts, page 16-11. - You will receive an alert if analysis of some file types is temporarily unavailable. See Ensuring That You Receive Alerts, page 16-11. - You will receive an alert if analysis of all supported file types is restored after a temporary outage. See Ensuring That You Receive Alerts, page 16-11. - Cisco AsyncOS for Email now includes a new message filter action (<code>skip-ampcheck</code>) that allows messages to bypass File Reputation Filtering and File Analysis configured on the system. See Bypass File Reputation Filtering and File Analysis System Actions, page 9-73.
Virtual Gateway Improvement	The number of Virtual Gateway addresses available on all Email Security appliance models is now 255. See Configuring Mail Gateways for all Hosted Domains Using Virtual Gateway™ Technology, page 24-59 .
Per-user spam notifications	You can specify which users receive spam notifications, based on LDAP groups.
Customizable end user notification page for URL filtering	You can customize the appearance of the end user notification page used for URL filtering and display your organization's branding such as logo, name of the organization, contact information, and so on. See Customizing the Appearance of End User Notification Page, page 15-6 .
Disk space management improvements	More than 2 TB of disk space is now available on physical hardware and on virtual appliances running ESXi 5.5 and VMFS 5. You can now allocate disk space on the appliance based on the functionality your organization uses (for spam and system quarantines, reporting and tracking data, etc.) Previous limits on quarantine sizes have been removed. For virtual appliances, you can use VMWare tools to increase the disk space available to Email Security appliance instances. For details, see Managing Disk Space, page 33-16.
Configurable SSL Settings in FIPS Mode	In FIPS mode, you can now configure the Cipher Suites in the SSL settings, using the <code>sslconfig</code> command in CLI. For more information, see <i>Cisco AsyncOS for Email CLI Reference Guide</i>. Note You cannot change server and client methods in FIPS mode.

Feature	Description
Configurable SSH Server Settings	You can now configure the following SSH server settings using the <code>sshconfig</code> command in CLI: • Public Key Authentication Algorithms • Cipher Algorithms • KEX Algorithms • MAC Methods • Minimum Server Key Size See Managing Secure Shell (SSH) Keys, page 32-28.
Encrypt Sensitive Data in FIPS Mode	In FIPS mode, you can now encrypt: • Critical security parameters in your appliance • Swap space in your appliance. This helps to prevent any unauthorized access or forensic attacks when the physical security of the appliance is compromised. Use the <code>fipsconfig</code> command in CLI to enable encryption of sensitive data in the appliance. See Encrypting Sensitive Data in FIPS Mode, page 27-3.
Encrypt Sensitive Data in Configuration Files	You can now encrypt the critical security parameters in the appliance configuration file while exporting, emailing, or displaying it. See Saving and Exporting the Current Configuration File, page 33-8.
Permanently Delete Sensitive Data in the Appliance	You can now permanently delete sensitive data (critical security parameters) in your appliance using one of the following commands in CLI: • <code>wipedata</code> • <code>diagnostic > reload</code> See <i>Cisco AsyncOS for Email CLI Reference Guide</i>.
More Secure AsyncOS Updates and Upgrades	For enhanced security, AsyncOS now uses a stronger hashing algorithm, SHA-384, to verify the received updates and upgrades.
Configurable CLI Session Timeout	You can now specify how long a user can be logged into the Email Security appliance's CLI before AsyncOS logs the user out due to inactivity. See Configuring the CLI Session Timeout, page 32-26. Note The CLI session timeout applies only to the connections using Secure Shell (SSH), SCP, and direct serial connection.
Enhanced Security for DKIM Signing Keys in FIPS Mode	For enhanced security, if encryption of sensitive data in the appliance is enabled in FIPS mode, • Private keys are not displayed in plain text while editing an existing signing key. See Edit an Existing Signing Key, page 20-11. • Signing keys are encrypted while exporting. See Exporting Signing Keys, page 20-11.
Enhanced Security for DSA Host Keys in FIPS Mode	For enhanced security, in FIPS mode, AsyncOS for Email uses a 2048-bit DSA host key.

Feature	Description
Enhanced Security for Demonstration Certificate	The demonstration certificate is updated to use keys of size 2048 bits and 1024 bits for FIPS mode and non-FIPS mode, respectively.
Enhanced Password Options	When creating user accounts or changing passwords, there is now an option to auto-generate a password that meets the configured requirements.
Welcome Banner	You can configure Cisco AsyncOS for Email to display a welcome banner after a user successfully logs into the appliance through SSH, Telnet, FTP, or web interface. You can use the welcome banner to display internal security information or best practice instructions for the appliance. See Displaying a Message After Login , page 32-27.
New authorization protocol for outgoing SMTP authentication	Outgoing SMTP authentication now supports the following additional authorization protocol: LOGIN.
Enhanced spam protection capabilities	Cisco AsyncOS now has enhanced capabilities to detect and protect against new spam campaigns, for example, snowshoe spam.
More flexibility for choosing users for an incoming or outgoing policy	Prior to this release, an incoming or outgoing policy matches if any of the specified values (sender, receiver domains, or LDAP group names) in the policy matches. Cisco AsyncOS 9.0 for Email provides you more flexibility for choosing users for an incoming or outgoing policy. You can set the policy to match if, • The message is from any sender, one or more of the specified senders, or none of the specified senders. • The message is sent to any recipient, one or more of the specified recipients, or all of the specified recipients and none of the specified recipients. Note From Cisco AsyncOS 9.0 for Email onwards, you must specify at least one sender and recipient. See Defining Senders and Recipients for Mail Policies, page 10-8.
Enhanced URL Defanging	Message and content filters for URL defanging now accounts for DNS spoofing and replaces a “.” (dot) in the URL with “[.]”. For example, after defanging, www.defangurl.com becomes BLOCKEDwww[.]defangurl[.]comBLOCKED.
Changed Behavior	
The disk_usage command has been deprecated	The disk_usage subcommand under diagnostics has been deprecated. To view and configure disk space quotas, use the diskquotaconfig command instead.
Opening a Support Case from the appliance	In order to open a support case from the appliance, you will need your CCOID and support contract number. Previously, this information was collected via other means. Also, in order to route cases more efficiently, the Technology and Sub-Technology options may differ from previous releases and may change at any time.

Feature	Description
Changes in Password Change Options	When you are enforcing a password change, you can choose whether the users must change the password during the next login or after a specified duration. If you are enforcing a password change after a specified duration, you can also set a grace period to reset the password after the password expires. See Force Users To Change Their Passwords, page 32-5 .
Changes in Local User Account and Password Settings	While configuring Local User Account and Password Settings, you can configure a grace period to reset the password after the password expires. See Configuring Restrictive User Account and Password Settings, page 32-17 .
Disk Space for Quarantines	You must now allocate disk space for quarantines using the System Administration > Disk Management menu.
New log for URL Filtering	URL filtering information will be posted to the following logs: - Mail Logs (<code>mail_logs</code>). Information related to the result of scanning a URL (action taken of a message depending on the URL) is posted to this log. - URL Filtering Logs (<code>web_client</code>). Information related to errors, timeouts, network issues, and so on while attempting the URL lookup are posted this log.
Stricter password rules	Stricter password rules are enforced immediately after running the System Setup Wizard.

Where to Find More Information

Cisco offers the following resources to learn more about your appliance:

- [Documentation, page 1-7](#)
- [Training, page 1-8](#)
- [Cisco Notification Service, page 1-8](#)
- [Knowledge Base, page 1-9](#)
- [Cisco Support Community, page 1-9](#)
- [Cisco Customer Support, page 1-9](#)
- [Third Party Contributors, page 1-9](#)
- [Cisco Welcomes Your Comments, page 1-10](#)
- [Registering for a Cisco Account, page 1-10](#)

Documentation

You can access the online help version of this user guide directly from the appliance GUI by clicking Help and Support in the upper-right corner.

The documentation set for the Cisco Email Security appliances includes the following documents and books:

- Release Notes
- The Quick Start Guide for Email Security appliances
- *Cisco AsyncOS for Email User Guide* (this book)
- *Cisco Content Security Virtual Appliance Installation Guide*
- *Cisco AsyncOS CLI Reference Guide*
- *Cisco AsyncOS API for Email - Getting Started Guide*

Documentation for all Cisco Content Security products is available from:

Documentation For Cisco Content Security Products	Location
Hardware and virtual appliances	See the applicable product in this table.
Cisco Email Security	http://www.cisco.com/c/en/us/support/security/content-security-management-appliance/tsd-products-support-series-home.html
Cisco Web Security	http://www.cisco.com/c/en/us/support/security/web-security-appliance/tsd-products-support-series-home.html
Cisco Content Security Management	http://www.cisco.com/c/en/us/support/security/email-security-appliance/tsd-products-support-series-home.html
CLI reference guide for Cisco Content Security appliances	http://www.cisco.com/c/en/us/support/security/email-security-appliance/products-command-reference-list.html
Cisco IronPort Encryption	http://www.cisco.com/c/en/us/support/security/email-encryption/tsd-products-support-series-home.html

Training

More information about training is available from:

- http://www.cisco.com/web/learning/le31/email_sec/index.html
- <http://www.cisco.com/web/learning/training-index.html>

Cisco Notification Service

Sign up to receive notifications relevant to your Cisco Content Security Appliances, such as Security Advisories, Field Notices, End of Sale and End of Support statements, and information about software updates and known issues.

You can specify options such as notification frequency and types of information to receive. You should sign up separately for notifications for each product that you use.

To sign up, visit <http://www.cisco.com/cisco/support/notifications.html>

A Cisco.com account is required. If you do not have one, see [Registering for a Cisco Account](#), page 1-10.

Knowledge Base

-
Step 1	Go to the main product page (http://www.cisco.com/c/en/us/support/security/email-security-appliance/tsd-products-support-series-home.html)
Step 2	Look for links with TechNotes in the name.
-

Cisco Support Community

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

Access the Cisco Support Community on the Customer Support Portal at the following URLs:

- For email security and associated management:
<https://supportforums.cisco.com/community/netpro/security/email>
- For web security and associated management:
<https://supportforums.cisco.com/community/netpro/security/web>

Cisco Customer Support

Use the following methods to obtain support:

U.S.: Call 1 (408) 526-7209 or Toll-free 1 (800) 553-2447

International: http://www.cisco.com/en/US/support/tsd_cisco_worldwide_contacts.html

Support Site: http://www.cisco.com/en/US/products/ps11169/serv_group_home.html

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

Third Party Contributors

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

The full text of these agreements can be found here:

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

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

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

Cisco Welcomes Your Comments

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

`contentsecuritydocs@cisco.com`

Please include the product name, release number, and document publication date in the subject of your message.

Registering for a Cisco Account

Access to many resources on Cisco.com requires a Cisco account.

If you do not have a Cisco.com User ID, you can register for one here:

<https://tools.cisco.com/RPF/register/register.do>

Related Topics

- [Cisco Notification Service, page 1-8](#)
- [Knowledge Base, page 1-9](#)

Cisco Email Security Appliance Overview

The Cisco AsyncOS™ operating system includes the following features:

- **Anti-Spam** at the gateway, through the unique, multi-layer approach of SenderBase Reputation Filters and Cisco Anti-Spam integration.
- **Anti-Virus** at the gateway with the Sophos and McAfee Anti-Virus scanning engines.
- **Outbreak Filters™**, Cisco's unique, preventive protection against new virus, scam, and phishing outbreaks that can quarantine dangerous messages until new updates are applied, reducing the window of vulnerability to new message threats.
- **Policy, Virus, and Outbreak Quarantines** provide a safe place to store suspect messages for evaluation by an administrator.
- **Spam Quarantine** either on-box or off, providing end user access to quarantined spam and suspected spam.
- **Email Authentication**. Cisco AsyncOS supports various forms of email authentication, including Sender Policy Framework (SPF), Sender ID Framework (SIDF), and DomainKeys Identified Mail (DKIM) verification of incoming mail, as well as DomainKeys and DKIM signing of outgoing mail.
- **Cisco Email Encryption**. You can encrypt outgoing mail to address HIPAA, GLBA and similar regulatory mandates. To do this, you configure an encryption policy on the Email Security appliance and use a local key server or hosted key service to encrypt the message.
- **Email Security Manager**, a single, comprehensive dashboard to manage all email security services and applications on the appliance. Email Security Manager can enforce email security based on user groups, allowing you to manage Cisco Reputation Filters, Outbreak Filters, Anti-Spam, Anti-Virus, and email content policies through distinct inbound and outbound policies.
- **On-box Quarantine areas** to hold messages that violate email policies. Quarantines seamlessly interact with the Outbreak Filters feature.

- **On-box message tracking.** AsyncOS for Email includes an on-box message tracking feature that makes it easy to find the status of messages that the Email Security appliance processes.
- **Mail Flow Monitoring** of all inbound and outbound email that provides complete visibility into all email traffic for your enterprise.
- **Access control** for inbound senders, based upon the sender's IP address, IP address range, or domain.
- Extensive **message filtering** technology allows you to enforce corporate policy and act on specific messages as they enter or leave your corporate infrastructure. Filter rules identify messages based on message or attachment content, information about the network, message envelope, message headers, or message body. Filter actions allow messages to be dropped, bounced, archived, blind carbon copied, or altered, or to generate notifications.
- **Message encryption via secure SMTP over Transport Layer Security** ensures messages traveling between your corporate infrastructure and other trusted hosts are encrypted.
- **Virtual Gateway™** technology allows the Email Security appliance to function as several email gateways within a single server, which allows you to partition email from different sources or campaigns to be sent over separate IP addresses. This ensures that deliverability issues affecting one IP address do not impact others.

AsyncOS for email supports RFC 2821-compliant Simple Mail Transfer Protocol (SMTP) to accept and deliver messages.

Most reporting, monitoring, and configuration commands are available through both the web-based GUI via HTTP or HTTPS. In addition, an interactive Command Line Interface (CLI) which you access from a Secure Shell (SSH), telnet, or direct serial connection is provided for the system.

You can also set up a Security Management appliance to consolidate reporting, tracking, and quarantine management for multiple Email Security appliances.

Related Topics

- [Supported Languages, page 1-11](#)

Supported Languages

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

- English
- French
- Spanish
- German
- Italian
- Korean
- Japanese
- Portuguese (Brazil)
- Chinese (traditional and simplified)
- Russian



Accessing the Appliance

- [Web-based Graphical User Interface \(GUI\), page 2-1](#)
- [Command Line Interface \(CLI\), page 2-3](#)

Web-based Graphical User Interface (GUI)

You can administer the appliance using both the web-based Graphical User Interface (GUI) and Command Line Interface (CLI). The GUI contains most of the functionality you need to configure and monitor the system. However, not all CLI commands are available in the GUI; some features are *only* available through the CLI.

- [Browser Requirements, page 2-1](#)
- [Accessing the GUI, page 2-1](#)

Browser Requirements

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

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

Do not use multiple browser windows or tabs simultaneously to make changes to the appliance. Do not use concurrent GUI and CLI sessions. Doing so will cause unexpected behavior and is not supported.

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

Accessing the GUI

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

<http://192.168.42.42>

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

Related Topics


- [Factory Default Username and Password, page 2-2](#)
- [Centralized Management, page 2-2](#)

Factory Default Username and Password

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

For example:

Figure 2-1 *The Login Screen*



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

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

For example:

`http://192.168.1.1` OR
`https://192.168.1.1` OR
`http://mail3.example.com` OR
`https://mail3.example.com`

**Note**

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

Related Topics

- [Adding Users, page 32-4](#)

Centralized Management

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

For more information, see [Administering a Cluster from the GUI, page 39-15](#).

Changing Configuration Settings

- [Configuration Changes, page 2-3](#)
- [Commit or Abandoning Changes, page 2-3](#)

Configuration Changes

You can make configuration changes while email operations proceed normally.

Commit or Abandoning Changes

You must explicitly save most configuration changes.

When changes are pending a commit, the Commit Changes button turns orange.

Figure 2-2 *The Commit Changes Button*



To clear or commit these changes, click **Commit Changes**.

Related Topics

- [Clearing Configuration Changes, page 2-8](#)

Command Line Interface (CLI)

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

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

Related Topics

- [Command Line Interface Conventions, page 2-3](#)
- [General Purpose CLI Commands, page 2-7](#)

Command Line Interface Conventions

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

- [Command Prompt, page 2-4](#)
- [Command Syntax, page 2-5](#)
- [Select Lists, page 2-5](#)
- [Yes/No Queries, page 2-5](#)

- [Subcommands, page 2-5](#)
- [History, page 2-6](#)
- [Command Completion, page 2-6](#)
- [Configuration Changes, page 2-6](#)

Command Prompt

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

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

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

```
(Cluster Americas) >
```

or

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

See [Centralized Management, page 2-2](#) for more information.

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

For example:

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

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

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

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

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

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

Command Syntax

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

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

Select Lists

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

For example:

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

Yes/No Queries

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

For example:

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

Subcommands

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

For example:

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

```
Currently configured interfaces:
```

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

```
Choose the operation you want to perform:
```

- NEW - Create a new interface.
- EDIT - Modify an interface.

```
- GROUPS - Define interface groups.
- DELETE - Remove an interface.

[]>
```

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

Escape

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

History

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

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

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

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

Command Completion

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

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

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

Configuration Changes

You can make configuration changes to Cisco AsyncOS while email operations proceed normally. Configuration changes will not take effect until you:

1. Issue the `commit` command at the command prompt.
2. Give the `commit` command the input required.
3. Receive confirmation of the `commit` procedure at the CLI.

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



Note Not all commands in AsyncOS require the `commit` command to be run. See the *Cisco AsyncOS CLI Reference Guide* for a summary of commands that require commit to be run before their changes take effect.

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

General Purpose CLI Commands

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

- [Committing Configuration Changes, page 2-7](#)
- [Clearing Configuration Changes, page 2-8](#)
- [Rolling Back Configuration Changes, page 2-8](#)
- [Quitting the Command Line Interface Session, page 2-9](#)
- [Seeking Help on the Command Line Interface, page 2-9](#)

Committing Configuration Changes

The `commit` command is critical to saving configuration changes to the appliance. Many configuration changes are not effective until you enter the `commit` command. (A few commands do not require you to use the `commit` command for changes to take effect. The `commit` command applies configuration changes made to Cisco AsyncOS since the last `commit` command or the last `clear` command was issued. You may include comments up to 255 characters. Changes are not verified as committed until you receive confirmation along with a timestamp.

Entering comments after the `commit` command is optional.

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

```
Please enter some comments describing your changes:
```

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

```
Do you want to save the current configuration for rollback? [Y]> n
```

```
Changes committed: Fri May 23 11:42:12 2014 GMT
```

**Note**

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

Clearing Configuration Changes

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

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

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

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

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

Rolling Back Configuration Changes

The `rollbackconfig` command lists the last ten committed configurations and allows you to select one to which you want to roll back.

Only administrators can use this command.

**Note**

This command does not work on clustered appliances. The appliance will not restore the configurations if you revert the appliance to an earlier version of AsyncOS.

```
mail.example.com> rollbackconfig
```

```
Previous Commits :
```

Committed On	User	Description

1. Wed Sep 19 22:03:10 2012	admin	Enabled anti-spam
2. Wed Sep 19 21:51:14 2012	admin	Updated envelope encry...
3. Wed Sep 19 18:50:41 2012	admin	

Enter the number of the config to revert to.

[> 1

Reverted to Wed Sep 19 18:50:41 2012 admin

Do you want to commit this configuration now? [N]> y

Committed the changes successfully

Quitting the Command Line Interface Session

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

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

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

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

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

Seeking Help on the Command Line Interface

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

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




Setup and Installation

- [Installation Planning, page 3-1](#)
- [Physically Connecting the Email Security Appliance to the Network, page 3-5](#)
- [Preparing for System Setup, page 3-8](#)
- [Using the System Setup Wizard, page 3-13](#)
- [Verifying Your Configuration and Next Steps, page 3-37](#)

Installation Planning

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

Review Information That Impacts Planning Decisions

- If you are configuring a virtual Email Security appliance, please see the *Cisco Content Security Virtual Appliance Installation Guide* before continuing with this chapter.
- If you are configuring an M-Series Cisco Content Security Management appliance, please see [Chapter 42, “Centralizing Services on a Cisco Content Security Management Appliance”](#).
- We recommend reviewing [Chapter 4, “Understanding the Email Pipeline”](#) before installing, as some features and functions may affect the placement of the appliance within your infrastructure.

Plan to Place the Email Security Appliance at the Perimeter of Your Network

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

The per-recipient reputation filtering, anti-spam, anti-virus, and Virus Outbreak Filter features (see [SenderBase Reputation Service, page 6-1](#), [IronPort Anti-Spam Filtering, page 13-3](#), [Sophos Anti-Virus Filtering, page 12-2](#), and [Outbreak Filters, page 14-1](#)) are designed to work with a *direct flow* of

messages from the Internet and from your internal network. You can configure the appliance for policy enforcement ([Overview of Defining Which Hosts Are Allowed to Connect, page 7-1](#)) for all email traffic to and from your enterprise.

Ensure that the Email Security 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 Email Security appliance will not be able to determine the sender’s IP address. The sender’s IP address is needed to identify and distinguish senders in the Mail Flow Monitor, to query the SenderBase Reputation Service for the sender’s SenderBase Reputation Score (SBRS), and to improve the efficacy of the Anti-Spam and Outbreak Filters features.

**Note**

If you cannot configure the appliance as the *first* machine receiving email from the Internet, you can still exercise some of the security services available on the appliance. For more information, see [Determining Sender IP Address In Deployments with Incoming Relays, page 13-15](#).

When you use the Email Security appliance as your SMTP gateway:

- The Mail Flow Monitor feature (see [Chapter 28, “Using Email Security Monitor”](#)) offers complete visibility into all email traffic for your enterprise from both internal and external senders.
- LDAP queries (see [Chapter 25, “LDAP Queries”](#)) for routing, aliasing, and masquerading can consolidate your directory infrastructure and provide for simpler updates.
- Familiar tools like alias tables (see [Creating Alias Tables, page 24-7](#)), domain-based routing ([The Domain Map Feature, page 24-28](#)), and masquerading ([Configuring Masquerading, page 24-16](#)) make the transition from Open-Source MTAs easier.

Register the Email Security Appliance in DNS

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

To register the 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 Email Security appliance as either a primary or backup MTA for your domain.

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

```
$ host -t mx example.com

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

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

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

Installation Scenarios

You can install your Email Security appliance into your existing network infrastructure in several ways. Most customers' network configurations are represented in the following scenarios. If your network configuration varies significantly and you would like assistance planning an installation, please contact Cisco Customer Support (see [Cisco Customer Support](#), page 1-9).

- [Configuration Overview](#), page 3-3
- [Incoming](#), page 3-3
- [Outgoing](#), page 3-3
- [Ethernet Interfaces](#), page 3-4
- [Advanced Configurations](#), page 3-4
- [Firewall Settings \(NAT, Ports\)](#), page 3-4

Configuration Overview

The following figure shows the typical placement of the Email Security appliance in an enterprise network environment:



In some scenarios, the Email Security appliance resides inside the network “DMZ,” in which case an additional firewall sits between the Email Security appliance and the groupware server.

The following network scenarios are described:

- Behind the Firewall: two listeners configuration ([Figure 3-1 on page 3-6](#))

Choose the configuration that best matches your infrastructure. Then proceed to the next section, [Preparing for System Setup](#), page 3-8.

Incoming

- Incoming mail is accepted for the local domains you specify.
- All other domains are rejected.
- External systems connect directly to the Email Security appliance to transmit email for the local domains, and the Email Security appliance relays the mail to the appropriate groupware servers (for example, Exchange™, Groupwise™, Domino™) via SMTP routes. (See [Routing Email for Local Domains](#), page 24-1.)

Outgoing

- Outgoing mail sent by internal users is routed by the groupware server to the Email Security appliance.
- The Email Security appliance accepts outbound email based on settings in the Host Access Table for the private listener. (For more information, see [Working with Listeners](#), page 5-2.)

Ethernet Interfaces

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

For more information about assigning multiple IP addresses to the available interfaces, see [Configuring Mail Gateways for all Hosted Domains Using Virtual Gateway™ Technology, page 24-59](#) and [Appendix B, “Assigning Network and IP Addresses”](#).

Hardware Ports

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

Ports	Type	C170	C370	C670	X1070	C380	C680
Management	Ethernet	0	1	1	1	1	1
Data	Ethernet	2*	3	3	3	3	3
Console	Serial	9-pin	9-pin	9-pin	9-pin	RJ-45	RJ-45
Remote Power Management (RPC)	Ethernet	N	N	N	N	Y	Y

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

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

Related Topics

- [Configuring Network Interfaces, page 3-17](#)
- [Accessing the Email Security appliance via a Serial Connection, page A-5](#)
- [Enabling Remote Power Management, page 33-29](#)

Advanced Configurations

In addition to the configurations shown in [Figure 3-1](#) and [Figure 3-2](#), you can also configure:

- Multiple Email Security appliances using the Centralized Management feature. See [Chapter 39, “Centralized Management Using Clusters.”](#)
- Redundancy at the network interface card level by “teaming” two of the Ethernet interfaces on Email Security appliances using the NIC Pairing feature. See [Chapter 37, “Advanced Network Configuration.”](#)

Firewall Settings (NAT, Ports)

SMTP and DNS services must have access to the Internet. Other services may also require open firewall ports. For details, see [Appendix D, “Firewall Information”](#).

Physically Connecting the Email Security Appliance to the Network

- [Configuration Scenarios, page 3-5](#)

Configuration Scenarios

The typical configuration scenario for the Email Security appliance is as follows:

- **Interfaces** - Only one of the three available Ethernet interfaces on the Email Security appliance is required for most network environments. However, you can configure two Ethernet interfaces and segregate your internal network from your external Internet network connection.
- **Public Listener (incoming email)** - The public listener receives connections from many external hosts and directs messages to a limited number of internal groupware servers.
 - Accepts connections from external mail hosts based on settings in the Host Access Table (HAT). By default, the HAT is configured to ACCEPT connections from all external mail hosts.
 - Accepts incoming mail only if it is addressed for the local domains specified in the Recipient Access Table (RAT). All other domains are rejected.
 - Relays mail to the appropriate internal groupware server, as defined by SMTP Routes.
- **Private Listener (outgoing email)** - The private listener receives connections from a limited number of internal groupware servers and directs messages to many external mail hosts.
 - Internal groupware servers are configured to route outgoing mail to the Cisco C- or X-Series appliance.
 - The Email Security 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.

Related Topics

- [Segregating Incoming and Outgoing Mail, page 3-5](#)


Segregating Incoming and Outgoing Mail

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

- 2 separate listeners on 2 logical IPv4 and 2 IPv6 addresses configured on *separate* physical interfaces
 - segregates incoming and outgoing traffic
 - you can assign an IPv4 and an IPv6 address to each listener
- 1 listener on 1 logical IPv4 address configured on one physical interface
 - combines both incoming and outgoing traffic
 - you can assign both an IPv4 and an IPv6 address to the listener

Configuration worksheets for both one and two listener configurations are included below (see [Gathering the Setup Information, page 3-11](#)). Most configuration scenarios are represented by one of the following three figures.

Figure 3-1 Behind the Firewall Scenario / 2 Listeners Configuration



Notes:

- 2 Listeners
- 2 IPv4 addresses
- 2 IPv6 addresses
- 1 or 2 Ethernet interfaces (only 1 interface shown)
- SMTP routes configured

Inbound Listener: "InboundMail" (public)

- IPv4 address: 1.2.3.4
- IPv6 address: 2001:0db8:85a3::8a2e:0370:7334
- Listener on the Data2 interface listens on port 25
- HAT (accept ALL)
- RAT (accept mail for local domains; reject ALL)


Outbound Listener: "OutboundMail" (private)

- IP address: 1.2.3.5
- IPv6 address: 2001:0db8:85a3::8a2e:0370:7335
- Listener on the Data2 interface listens on port 25
- HAT (relay for local domains; reject ALL)

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

SMTP routes direct mail to proper groupware server

Firewall ports opened for appropriate services to and from the Email Security appliance

Figure 3-2 One Listener Configuration**Notes:**

- 1 Listener
- 1 IP addresses
- 1 Ethernet interface
- SMTP routes configured

Inbound Listener: "InboundMail" (public)

- IP address: 1.2.3.4
- Listener on the Data2 interface listens on port 25
- HAT (accept ALL) includes entries for Groupware servers in RELAYLIST
- RAT (accept mail for local domains; reject ALL)

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

SMTP routes direct mail to proper groupware server

Firewall ports opened for appropriate services to and from the appliance

Preparing for System Setup

- [Determine Method for Connecting to the Appliance, page 3-8](#)
- [Determining Network and IP Address Assignments, page 3-9](#)
- [Gathering the Setup Information, page 3-11](#)

	Do This	More Information
Step 1	Determine how you will connect to the appliance.	See Determine Method for Connecting to the Appliance, page 3-8
Step 2	Determine network and IP address assignments. If you have already cabled your appliance to your network, ensure that the default IP address for the Email Security appliance does not conflict with other IP addresses on your network.	Determining Network and IP Address Assignments, page 3-9 Determine Method for Connecting to the Appliance, page 3-8
Step 3	Gather information about your system setup.	See Gathering the Setup Information, page 3-11 .
Step 4	Review the latest product release notes for your appliance.	Release notes are available from the link in Documentation, page 1-7 .
Step 5	Unpack the appliance, physically install it in a rack, and turn it on.	See <i>Quickstart Guide</i> for your appliance. This guide is available from the link in Documentation, page 1-7 .
Step 6	Access the appliance using the web interface or the command line interface (CLI)	• Launch a web browser and enter the IP address of the appliance. or • See Running the Command Line Interface (CLI) System Setup Wizard, page 3-24
Step 7	If you are setting up a virtual Email Security appliance, load your virtual appliance license.	Use the <code>loadlicense</code> command. For more information, see the <i>Cisco Content Security Virtual Appliance Installation Guide</i> available from the link in Documentation, page 1-7 .
Step 8	Configure basic settings for your system.	See Using the System Setup Wizard, page 3-13

Determine Method for Connecting to the Appliance

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

Related Topics

- [Connecting to the Appliance, page 3-9](#)

Connecting to the Appliance

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

Table 3-1 **Options for Connecting to the Appliance**

Ethernet	An Ethernet connection between a PC and the network and between the network and the Management port. The IPv4 address that has been assigned to the Management port by the factory is 192.168.42.42. This is the easiest way to connect if it works with your network configuration.
Serial	A serial communications connection between the PC and the Serial Console port. If you cannot use the Ethernet method, a straight serial-to-serial connection between the computer and the appliance will work until alternate network settings can be applied to the Management port. For pinout information, see Accessing the Email Security appliance via a Serial Connection, page A-5. The communications settings for the serial port are: Bits per second: 9600 Data bits: 8 Parity: None Stop bits: 1 Flow control: Hardware



Note

Keep in mind that the initial connection method is not final. This process applies only for the initial configuration. You can change network settings at a later time to allow different connection methods. (See [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) for more information.) You can also create multiple user accounts with differing administrative privileges to access the appliance. (For more information, see [Adding Users, page 32-4](#).)

Determining Network and IP Address Assignments

You can use both IPv4 and IPv6 addresses.

- [Default IP Addresses for Management and Data Ports, page 3-9](#)
- [Choosing Network Connections to Receive and Deliver Email, page 3-10](#)
- [Binding Logical IP Addresses to Physical Ethernet Ports, page 3-10](#)
- [Choosing Network Settings for Your Connections, page 3-10](#)

Default IP Addresses for Management and Data Ports

The IP address that is pre-configured on the Management port (the Data 1 port on C170 appliances) is 192.168.42.42.

Choosing Network Connections to Receive and Deliver Email

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

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

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

Binding Logical IP Addresses to Physical Ethernet Ports

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

- 2 separate listeners on 2 logical IPv4 and 2 IPv6 addresses configured on *separate* physical interfaces
 - segregates incoming and outgoing traffic
 - you can assign an IPv4 and an IPv6 address to each listener
- 1 listener on 1 logical IPv4 address configured on one physical interface
 - combines both incoming and outgoing traffic
 - you can assign both an IPv4 and an IPv6 address to the listener

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

Choosing Network Settings for Your Connections

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

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

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

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

See [Appendix B, "Assigning Network and IP Addresses"](#) for more information.



Note

If you are running a firewall on your network between the Internet and the Email Security appliance, it may be necessary to open specific ports for the appliance to work properly. See [Appendix D, "Firewall Information"](#) for more information.

Gathering the Setup Information

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

See [Appendix B, “Assigning Network and IP Addresses”](#) for more detailed information on network and IP addresses. See [Chapter 42, “Centralizing Services on a Cisco Content Security Management Appliance”](#) if you are configuring a Cisco Content Security Management appliance.

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

System Settings		
Default System Hostname:		
Email System Alerts To:		
Deliver Scheduled Reports To:		
Time Zone Information:		
NTP Server:		
Admin Password:		
SenderBase Network Participation:	Enable / Disable	
AutoSupport:	Enable / Disable	
Network Integration		
Gateway:		
DNS (Internet or Specify Own):		
Interfaces		
Data 1 Port		
IPv4 Address / Netmask:		
IPv6 Address / Prefix:		
Fully Qualified Hostname:		
Accept Incoming Mail:	Domain	Destination
Relay Outgoing Mail:	System	
Data 2 Port		
IPv4 Address / Netmask:		
IPv6 Address / Prefix:		
Fully Qualified Hostname:		
Accept Incoming Mail:	Domain	Destination
Relay Outgoing Mail:	System	
Management Port		
IP Address:		
Network Mask:		

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

IPv6 Address:		
Prefix:		
Fully Qualified Hostname:		
Accept Incoming Mail:	Domain	Destination
Relay Outgoing Mail:	System	
Message Security		
SenderBase Reputation Filtering:	Enable / Disable	
Anti-Spam Scanning Engine	None / IronPort	
McAfee Anti-Virus Scanning Engine	Enable / Disable	
Sophos Anti-Virus Scanning Engine	Enable / Disable	
Outbreak Filters	Enable / Disable	

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

System Settings		
Default System Hostname:		
Email System Alerts To:		
Deliver Scheduled Reports To:		
Time Zone:		
NTP Server:		
Admin Password:		
SenderBase Network Participation:	Enable / Disable	
AutoSupport:	Enable / Disable	
Network Integration		
Gateway:		
DNS (Internet or Specify Own):		
Interfaces		
Data2 Port		
IPv4 Address / Netmask:		
IPv6 Address / Prefix:		
Fully Qualified Hostname:		
Accept Incoming Mail:	Domain	Destination
Relay Outgoing Mail:	System	
Data1 Port		
IPv4 Address / Netmask:		

Table 3-3 **System Setup Worksheet: 1 Listener for All Email Traffic (continued)**

IPv6 Address / Prefix:	
Fully Qualified Hostname:	
Message Security	
SenderBase Reputation Filtering:	Enable / Disable
Anti-Spam Scanning Engine	None / IronPort
McAfee Anti-Virus Scanning Engine	Enable / Disable
Sophos Anti-Virus Scanning Engine	Enable / Disable
Outbreak Filters	Enable / Disable

Using the System Setup Wizard

- [Accessing the Web-Based Graphical User Interface \(GUI\), page 3-14](#)
- [Defining Basic Configuration Using the Web-Based System Setup Wizard, page 3-14](#)
- [Setting up the Connection to Active Directory, page 3-22](#)
- [Proceeding to the Next Steps, page 3-23](#)
- [Accessing the Command Line Interface \(CLI\), page 3-23](#)
- [Running the Command Line Interface \(CLI\) System Setup Wizard, page 3-24](#)
- [Configuring your system as an Enterprise Gateway, page 3-37](#)

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

You can run the System Setup Wizard using a browser or the command line interface (CLI). For more information, see [Accessing the Web-Based Graphical User Interface \(GUI\), page 3-14](#) or [Running the Command Line Interface \(CLI\) System Setup Wizard, page 3-24](#)

Before you begin, complete the prerequisites at [Preparing for System Setup, page 3-8](#).



Caution

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



Caution

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



Caution

The Email Security appliance ships with a default IP address of 192.168.42.42 on the Management port of all hardware except C170 appliances, which use the Data 1 port instead. Before connecting the appliance to your network, ensure that no other device's IP address conflicts with this factory default setting. If you are configuring a Cisco Content Security Management appliance, please see [Centralizing Services on a Cisco Content Security Management Appliance, page 42-1](#).


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

Accessing the Web-Based Graphical User Interface (GUI)

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

The login screen is displayed:

Figure 3-3 *Logging in to the Appliance*
Welcome



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

Related Topics

- [Factory Default Username and Password, page 3-14](#)

Factory Default Username and Password

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



Note

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

Defining Basic Configuration Using the Web-Based System Setup Wizard

Procedure

-
- Step 1** Launch the System Setup Wizard
- Log in to the graphical user interface as described in [Accessing the Web-Based Graphical User Interface \(GUI\), page 3-14](#).
 - On brand new (not upgraded from previous releases of AsyncOS) systems, your browser will automatically be redirected to the System Setup Wizard.
 - Otherwise, on the System Administration tab, click System Setup Wizard in the list of links on the left.
- Step 2** Start. See [Step 1: Start](#).
- Read and accept the license agreement
- Step 3** System. See [Step 2: System](#).

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

Step 4 Network. See [Step 3: Network](#).

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

Step 5 Security. See [Step 4: Security](#).

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

Step 6 Review. See [Step 5: Review](#).

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

Step 7 Commit the changes you have made.

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

Step 1: Start

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

You can also view the text of the agreement here:

<https://support.ironport.com/license/eula.html>

Step 2: System

- [Setting the Hostname, page 3-16](#)
- [Configuring System Alerts, page 3-16](#)
- [Configuring Report Delivery, page 3-16](#)
- [Setting the Time, page 3-16](#)
- [Setting the Password, page 3-16](#)
- [Participating in SenderBase Network, page 3-16](#)

- [Enabling AutoSupport, page 3-17](#)

Setting the Hostname

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

Configuring System Alerts

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

You must add at least one email address that receives System Alerts. Enter a single email address, or separate multiple addresses with commas. The email recipients initially receive all types of alerts at all levels, except for Directory Harvest Attack Prevention alerts. You can add more granularity to the alert configuration later. For more information, see [Alerts, page 33-34](#).

Configuring Report Delivery

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

Setting the Time

Set the time zone on the Email Security appliance so that timestamps in message headers and log files are correct. Use the drop-down menus to locate your time zone or to define the time zone via GMT offset (see [Selecting a GMT Offset, page 33-59](#) for more information).

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

Setting the Password

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

Participating in SenderBase Network

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

If you agree to participate in the SenderBase Network, Cisco will collect aggregated email traffic statistics about your organization. This includes only summary data on message attributes and information on how different types of messages were handled by Email Security appliances. For example, Cisco does not collect the message body or the message subject. Personally identifiable information or information that identifies your organization will be kept confidential. To learn more about SenderBase, including examples of the data collected, follow the **Click here for more information about what data is being shared...** link (see [Frequently Asked Questions, page 35-2](#)).

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

See [Chapter 35, “SenderBase Network Participation”](#) for more information.

Enabling AutoSupport

The AutoSupport feature (enabled by default) keeps the Cisco Customer Support team aware of issues with your appliance so that we can provide better support to you. (For more information, see [AutoSupport, page 33-34](#).)

Click **Next** to continue.

Step 3: Network

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

- [Configuring DNS and Default Gateway, page 3-17](#)
- [Configuring Network Interfaces, page 3-17](#)
- [Accepting Mail, page 3-18](#)
- [Relaying Mail \(Optional\), page 3-19](#)
- [C170 Installations, page 3-20](#)

Configuring DNS and Default Gateway

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

Next, configure the DNS (Domain Name Service) settings. Cisco AsyncOS contains a high-performance internal DNS resolver/cache that can query the Internet’s root servers directly, or the system can use DNS servers you specify. If you choose to use your own servers, you will need to supply the IP address and hostname of each DNS server. You can enter up to four DNS servers via the System Setup Wizard. Please note that DNS servers you enter will have an initial priority of 0. For more information, see [Configuring Domain Name System \(DNS\) Settings, page 33-54](#).



Note

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

Configuring Network Interfaces

Your Email Security appliance has network interfaces that are associated with the physical Ethernet ports on the machine.

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

Each interface can be configured to accept mail (incoming), relay email (outgoing), or appliance management. During setup, you are limited to one of each. On most appliances, you would typically use one interface for incoming, one for outgoing, and one for appliance management. On the C170 appliances, you would typically use one interface for both incoming and outgoing mail, and the other interface for management.

You must configure one interface to receive email.

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

For C370, C670, X1070, C380, C680 appliances: Cisco recommends using one of the physical Ethernet ports to connect directly to the Internet for the purposes of receiving inbound email through public listeners, and using another physical Ethernet port to connect directly to your internal network for the purposes of relaying outbound email through private listeners.

For C170 appliances: Typically, the System Setup Wizard will configure only one physical Ethernet port with one listener for both receiving inbound email and relaying outbound email.

See [Binding Logical IP Addresses to Physical Ethernet Ports, page 3-10](#).

The following information is required:

- The **IP address** assigned by your network administrator. This can be an IPv4 address, an IPv6 address, or both.
- For IPv4 addresses: the **netmask** of the interface. AsyncOS only accepts a netmask in CIDR format. For example, /24 for the 255.255.255.0 subnet.
For IPv6 addresses: the **prefix** in CIDR format. For example /64 for a 64-bit prefix.
- (optional) A fully-qualified hostname for the IP address.



Note

IP addresses within the same subnet cannot be configured on separate physical Ethernet interfaces. See [Appendix B, “Assigning Network and IP Addresses”](#) for more detailed information on Network and IP Address configuration.

Accepting Mail

When configuring your interfaces to accept mail, you define:

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

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

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

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

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

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

**Note**

Configuring SMTP Routes in this step is optional. If no SMTP routes are defined, the system will use DNS to lookup and determine the delivery host for the incoming mail received by the listener. (See [Routing Email for Local Domains, page 24-1.](#))

You must add at least one domain to the Recipient Access Table. Enter a domain —`example.com`, for example. To ensure that mail destined for any subdomain of `example.net` will match in the Recipient Access Table, enter `.example.net` as well as the domain name. For more information, see [Defining Recipient Addresses, page 8-4.](#)

Relaying Mail (Optional)

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

These are entries in the RELAYLIST of the Host Access Table for a listener. See [Sender Group Syntax, page 7-4](#) for more information.

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

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

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

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

**Note**

The following example pertains to the following appliances only: C370, C670, X1070, C380, C680. For C170 appliances, the Data 2 interface is typically configured for both incoming and outgoing mail while the Data 1 interface is used for appliance management (see [C170 Installations, page 3-20.](#))


Figure 3-4 Network Interfaces: 2 Interfaces in Addition to Management (Segregated Traffic)

☒ Enable Data 1 Interface	
<i>This interface is typically configured to accept mail.</i>	
IPv4 Address / Netmask:	1.1.1.1/24
IPv6 Address / Prefix:	2001:db8:1::4/64
Fully Qualified Hostname:	Fully qualified hostname for this appliance
Accept Incoming Mail:	☐ Accept mail on this interface
Relay Outgoing Mail:	☐ Relay mail on this interface
☒ Enable Data 2 Interface	
<i>This interface is typically configured to relay mail.</i>	
IPv4 Address / Netmask:	1.1.1.2/24
IPv6 Address / Prefix:	2001:db8:1::4/64
Fully Qualified Hostname:	Fully qualified hostname for this appliance
Accept Incoming Mail:	☐ Accept mail on this interface
Relay Outgoing Mail:	☐ Relay mail on this interface
☒ Enable Management Interface	
<i>This interface is typically configured for system administration.</i>	
IPv4 Address / Netmask:	1.1.1.2/24
IPv6 Address / Prefix:	2001:db8:1::4/64
Fully Qualified Hostname:	mail.example.com Fully qualified hostname for this appliance
Accept Incoming Mail:	☐ Accept mail on this interface
Relay Outgoing Mail:	☐ Relay mail on this interface

C170 Installations

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

Figure 3-5 Network Interfaces: 1 IP Address for Incoming and Outgoing (Nonsegregated) Traffic

Interfaces	
<i>You must set up at least 1 interface and 1 interface must be configured to accept mail from the Internet.</i>	
	
☒ Enable Data 2 Interface	
<i>This interface is typically used to accept and relay mail.</i>	
IP Address:	192.168.1.1
Network Mask:	255.255.255.0
Fully Qualified Hostname:	mail3.example.com Fully qualified hostname for this appliance
Accept Incoming Mail:	☒ Accept mail on this interface
Relay Outgoing Mail:	☒ Relay mail on this interface
Domain	example.com
Destination	exchange.example.com
example: company.com i.e. An Exchange or Notes server	
System	exchange.example.com
example: company.com	
☒ Enable Data 1 Interface	
<i>This interface is typically used for system administration. (You are currently connected to this interface.)</i>	
IP Address:	192.168.42.42
Network Mask:	255.255.255.0
Fully Qualified Hostname:	mail.example.com Fully qualified hostname for this appliance
Accept Incoming Mail:	☐ Accept mail on this interface
Relay Outgoing Mail:	☐ Relay mail on this interface

Click **Next** to continue.

Step 4: Security

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

- [Enabling SenderBase Reputation Filtering, page 3-21](#)
- [Enabling Anti-Spam Scanning, page 3-21](#)
- [Enabling Anti-Virus Scanning, page 3-21](#)
- [Enabling Advanced Malware Protection \(File Reputation and Analysis Services\), page 3-21](#)
- [Enabling Outbreak Filters, page 3-22](#)

Enabling SenderBase Reputation Filtering

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

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

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

Enabling Anti-Spam Scanning

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

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

See [Chapter 13, “Anti-Spam”](#) for all of the Anti-Spam configuration options available on the appliance. See [Policy, Virus, and Outbreak Quarantines, page 30-1](#).

Enabling Anti-Virus Scanning

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

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

See [Chapter 12, “Anti-Virus”](#) for all of the anti-virus configuration options available on the appliance.

Enabling Advanced Malware Protection (File Reputation and Analysis Services)

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

For more information, see [Chapter 16, “File Reputation Filtering and File Analysis.”](#)

Enabling Outbreak Filters

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

See [Chapter 14, “Outbreak Filters”](#) for more information.

Click **Next** to continue.

Step 5: Review

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

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

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

Your appliance is now ready to send email.



Note

Clicking **Install** will cause the connection to the current URL (<http://192.168.42.42>) to be lost if you changed the IP address of the interface you used to connect to the appliance (the Management interface on C370, C670, X1070, C380, C680 appliances, or the Data 1 interface on C170 appliances) from the default. However, your browser will be redirected to the new IP address.

Once System Setup is complete, several alert messages are sent. See [Immediate Alerts, page 3-37](#) for more information.

Setting up the Connection to Active Directory

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

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

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

Procedure

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

- Step 2** Enter the host name for the Active Directory server.
- Step 3** Enter a username and password for the authentication request.
- Step 4** Click **Next** to continue.
- The Active Directory Wizard tests the connection to the Active Directory server. If successful, the Test Directory Settings page is displayed.
- Step 5** Test the directory settings by entering an email address that you know exists in the Active Directory and clicking **Test**. The results appear in the connection status field.
- Step 6** Click **Done**.
-

Proceeding to the Next Steps

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

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

Accessing the Command Line Interface (CLI)

Access to the CLI varies depending on the management connection method you chose in [Connecting to the Appliance, page 3-9](#). The factory default username and password are listed next. Initially, only the admin user account has access to the CLI. You can add other users with differing levels of permission after you have accessed the command line interface for the first time via the admin account. (For information about adding users, see [Adding Users, page 32-4](#).) The System Setup Wizard asks you to change the password for the admin account. The password for the admin account can also be reset directly at any time using the password command.

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

To connect via a Serial connection: Start a terminal session with the communication port on your personal computer that the serial cable is connected to. Use the settings for serial port outlined in [Connecting to the Appliance, page 3-9](#). Enter the username and password below.

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

Related Topics

- [Factory Default Username and Password, page 3-23](#)

Factory Default Username and Password

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

For example:

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

Running the Command Line Interface (CLI) System Setup Wizard

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

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

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

```
IronPort> systemsetup
```

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

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

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

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



Note

The remainder of the system setup steps are described below. Examples of the CLI System Setup Wizard dialogue will only be included for sections that deviate from the GUI System Setup Wizard described above in [Defining Basic Configuration Using the Web-Based System Setup Wizard, page 3-14](#).

Related Topics

- [Change the Admin Password, page 3-25](#)
- [Accept the License Agreement, page 3-25](#)
- [Set the Hostname, page 3-25](#)
- [Assign and Configure Logical IP Interface\(s\), page 3-25](#)
- [Specify the Default Gateway, page 3-26](#)
- [Enable the Web Interface, page 3-26](#)
- [Configure the DNS Settings, page 3-26](#)

- [Create a Listener, page 3-27](#)
- [Enable Anti-Spam, page 3-34](#)
- [Select a Default Anti-Spam Scanning Engine, page 3-34](#)
- [Enable the Spam Quarantine, page 3-34](#)
- [Enable Anti-Virus Scanning, page 3-34](#)
- [Enable Outbreak FiltersOutbreak Filters and SenderBase Email Traffic Monitoring Network, page 3-35](#)
- [Configure the Alert Settings and AutoSupport, page 3-35](#)
- [Configure Scheduled Reporting, page 3-35](#)
- [Configure Time Settings, page 3-36](#)
- [Commit Changes, page 3-36](#)
- [Test the Configuration, page 3-36](#)
- [Immediate Alerts, page 3-37](#)

Change the Admin Password

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

Accept the License Agreement

Read and accept the software license agreement that is displayed.

Set the Hostname

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

Assign and Configure Logical IP Interface(s)

The next step assigns and configures a logical IP interface on the physical Ethernet interface named Management (on C370, C670, X1070, C380, C680 appliances) or Data 1 (on C170 appliances), and then prompts you to configure a logical IP interface on any other physical Ethernet interfaces available on the appliance.

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

For C370, C670, X1070, C380, C680 appliances: Cisco recommends using one of the physical Ethernet ports to connect directly to the Internet for the purposes of receiving inbound email through public listeners, and using another physical Ethernet port to connect directly to your internal network for the purposes of relaying outbound email through private listeners.

For C170 appliances: By default, the `systemsetup` command will configure only one physical Ethernet port with one listener for receiving inbound email and relaying outbound email.

**Note**

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

The following information is required:

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

**Note**

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

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

**Note**

IP addresses within the same subnet cannot be configured on separate physical Ethernet interfaces. See [Appendix B, “Assigning Network and IP Addresses”](#) for more detailed information on Network and IP Address configuration.

**Note**

For C170 appliances, the Data 2 interface is configured first.

Specify the Default Gateway

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

Enable the Web Interface

In the next portion of the `systemsetup` command, you enable the web interface for the appliance (for the Management Ethernet interface). You can also choose to run the web interface over secure HTTP (`https`). If you choose to use HTTPS, the system will use a demonstration certificate until you upload your own certificate. For more information, see [Enabling a Certificate for HTTPS, page 23-18](#).

Configure the DNS Settings

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

Create a Listener

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

For C370, C670, X1070, C380, C680 appliances: By default, the `systemsetup` command configures two listeners — one public and one private. (For more information on the types of listeners available, see [Configuring the Gateway to Receive Email, page 5-1](#).)

For C170 appliances: By default, the `systemsetup` command configures one public listener for both receiving mail from the Internet and for relaying email from your internal network. See [Listener Example for C170 Appliances, page 3-31](#).

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

- A **name** (nickname) created by you to refer to the listener later. For example, the listener that accepts email from your internal systems to be delivered to the Internet may be called OutboundMail.
- One of the IP interfaces (that you created earlier in the `systemsetup` command) on which to receive email.
- The name of the machine(s) to which you want to route email (public listeners only). (This is the first `smtproutes` entry. See [Routing Email for Local Domains, page 24-1](#).)
- Whether or not to enable filtering based on SenderBase Reputation Scores (SBRS) for public listeners. If enabled, you are also prompted to select between Conservative, Moderate, or Aggressive settings.
- Rate-limiting per host: the maximum number of recipients per hour you are willing to receive from a remote host (public listeners only).
- The recipient domains or specific addresses you want to accept email for (public listeners) or the systems allowed to relay email through the appliance (private listeners). (These are the first Recipient Access Table and Host Access Table entries for a listener. See [Sender Group Syntax, page 7-4](#) and [Adding Domains and Users For Which to Accept Messages, page 8-3](#) for more information.)

Related Topics

- [Public Listener, page 3-27](#)
- [Private Listener, page 3-30](#)
- [Listener Example for C170 Appliances, page 3-31](#)

Public Listener



Note

The following examples of creating a public and private listener apply to C370, C670, X1070, C380, C680 appliances only. For C170 appliances, skip to the next section, [Listener Example for C170 Appliances, page 3-31](#).

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

**Note**

The value you enter for maximum recipients per hour you are willing to receive from a remote host is a completely arbitrary value, one that is usually relative to the size of the enterprise for which you are administering email. For example, a sender who sends 200 messages in an hour might be considered a “spammer” (sender of unsolicited bulk email), but if you are configuring the Email Security appliance to handle all email for a 10,000 person company, 200 messages per hour from a remote host may be a reasonable value. Conversely, in a 50-person company, someone sending 200 messages in an hour to you may be an obvious spammer. You must choose an appropriate value when you enable rate-limiting on a public listener (throttle) inbound email for your enterprise. For more information on Default Host Access policies, see [Sender Group Syntax, page 7-4](#).

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

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

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

```
[> InboundMail
```

Please choose an IP interface for this Listener.

1. Management (192.168.42.42/24: mail3.example.com)
2. PrivateNet (192.168.1.1/24: mail3.example.com)
3. PublicNet (192.168.2.1/24: mail3.example.com)

```
[1]> 3
```

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

Hostnames such as "example.com" are allowed.

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

Usernames such as "postmaster@" are allowed.

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

Separate multiple addresses with commas.

```
[> example.com
```

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

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

```
[> exchange.example.com
```

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

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

```
[> 4500
```

```
Default Policy Parameters
```

```
=====
```

```
Maximum Message Size: 100M
```

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

```
Maximum Number Of Messages Per Connection: 1,000
```

```
Maximum Number Of Recipients Per Message: 1,000
```

```
Maximum Number Of Recipients Per Hour: 4,500
```

```
Maximum Recipients Per Hour SMTP Response:
```

```
452 Too many recipients received this hour
```

```
Use SenderBase for Flow Control: Yes
```

```
Virus Detection Enabled: Yes
```

```
Allow TLS Connections: No
```

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

```
Listener InboundMail created.
```

```
Defaults have been set for a Public listener.
```

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

Private Listener

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

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

Note that the default values for a private listener differ from the public listener created earlier. For more information, see [Working with Listeners, page 5-2](#).

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

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

[]> **OutboundMail**

Please choose an IP interface for this Listener.

1. Management (192.168.42.42/24: mail3.example.com)
2. PrivateNet (192.168.1.1/24: mail3.example.com)
3. PublicNet (192.168.2.1/24: mail3.example.com)

[1]> **2**

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

Hostnames such as "example.com" are allowed.

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

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

Separate multiple entries with commas.

[]> **.example.com**

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

[N]> **n**


```

Default Policy Parameters

=====

Maximum Message Size: 100M

Maximum Number Of Connections From A Single IP: 600

Maximum Number Of Messages Per Connection: 10,000

Maximum Number Of Recipients Per Message: 100,000

Maximum Number Of Recipients Per Hour: Disabled

Use SenderBase for Flow Control: No

Virus Detection Enabled: Yes

Allow TLS Connections: No

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

Listener OutboundMail created.

Defaults have been set for a Private listener.

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

*****

```

Listener Example for C170 Appliances



Note

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

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

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



Note

The value you enter for maximum recipients per hour you are willing to receive from a remote host is a completely arbitrary value, one that is usually relative to the size of the enterprise for which you are administering email. For example, a sender who sends 200 messages in an hour might be considered a “spammer” (sender of unsolicited bulk email), but if you are configuring the appliance to handle all email for a 10,000 person company, 200 messages per hour from a remote host may be a reasonable value. Conversely, in a 50-person company, someone sending 200 messages in an hour to you may be an

obvious spammer. You must choose an appropriate value when you enable rate-limiting on a public listener (throttle) inbound email for your enterprise. For more information on Default Host Access policies, see [Sender Group Syntax, page 7-4](#).

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

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

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

```
[> MailInterface
```

Please choose an IP interface for this Listener.

1. MailNet (10.1.1.1/24: mail3.example.com)
2. Management (192.168.42.42/24: mail3.example.com)

```
[1]> 1
```

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

Hostnames such as "example.com" are allowed.

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

Usernames such as "postmaster@" are allowed.

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

Separate multiple addresses with commas.

```
[> example.com
```

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

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

```
[> exchange.example.com
```

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

Hostnames such as "example.com" are allowed.

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

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

Separate multiple entries with commas.

```
[> .example.com
```

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

```
[Y]> y
```

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

```
[> 450
```

Default Policy Parameters

=====

Maximum Message Size: 10M

Maximum Number Of Connections From A Single IP: 50

Maximum Number Of Messages Per Connection: 100

Maximum Number Of Recipients Per Message: 100

Maximum Number Of Recipients Per Hour: 450

Maximum Recipients Per Hour SMTP Response:

452 Too many recipients received this hour

Use SenderBase for Flow Control: Yes

Spam Detection Enabled: Yes

Virus Detection Enabled: Yes

Allow TLS Connections: No

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

Listener MailInterface created.

Defaults have been set for a Public listener.

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

**Note**

Because the `systemsetup` command only configures one listener for both inbound and outbound mail for C170 appliances, all outgoing mail will be calculated in the Mail Flow Monitor feature (which is normally used for inbound messages). See [Chapter 28, “Using Email Security Monitor.”](#)

Enable Anti-Spam

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

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

**Note**

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

See [Chapter 13, “Anti-Spam”](#) for all of the Anti-Spam configuration options available on the appliance.

Select a Default Anti-Spam Scanning Engine

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

Enable the Spam Quarantine

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

See [Setting Up the Local Spam Quarantine, page 31-2](#).

Enable Anti-Virus Scanning

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

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

See [Chapter 12, “Anti-Virus”](#) for the anti-virus configuration options available on the appliance.

Enable Outbreak Filters and SenderBase Email Traffic Monitoring Network

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

Related Topics

- [Outbreak Filters, page 3-35](#)
- [SenderBase Participation, page 3-35](#)

Outbreak Filters

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

If you choose to enable Outbreak Filters, enter a threshold value and whether you would like to receive Outbreak Filters alerts. For more information about Outbreak Filters and threshold values, see [Outbreak Filters, page 14-1](#).

SenderBase Participation

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

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

See [Chapter 35, “SenderBase Network Participation”](#) for more information.

Configure the Alert Settings and AutoSupport

Cisco AsyncOS sends alert messages to a user via email if there is a system error that requires the user’s intervention. Add at least one email address that receives system alerts. Separate multiple addresses with commas. The email addresses that you enter initially receive all types of alerts at all levels, except for Directory Harvest Attack Prevention alerts. You can add more granularity to the alert configuration later using the `alertconfig` command in the CLI or the System Administration > Alerts page in the GUI. For more information, see [Alerts, page 33-34](#).

The AutoSupport feature keeps the Cisco Customer Support team aware of issues with your appliance so that Cisco can provide industry-leading support to you. Answer “Yes” to send Cisco support alerts and weekly status updates. (For more information, see [AutoSupport, page 33-34](#).)

Configure Scheduled Reporting

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

Configure Time Settings

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

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

Commit Changes

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

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

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

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

The appliance is now ready to send email.

Test the Configuration

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

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

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

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

```
[> user@example.com
```

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

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

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

Immediate Alerts

The Email Security appliance uses feature keys to enable features. The first time you create a listener in the `systemsetup` command, enable Anti-Spam, enable Sophos or McAfee Anti-Virus, or enable Outbreak Filters, an alert is generated and sent to the addresses you specified in [Step 2: System, page 3-15](#).

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

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

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

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

For information on enabling a feature beyond the 30-day evaluation period, contact your Cisco sales representative. You can see how much time remains on a key via the System Administration > Feature Keys page or by issuing the `featurekey` command. (For more information, see [Feature Keys, page 33-5](#).)

Configuring your system as an Enterprise Gateway

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

Verifying Your Configuration and Next Steps

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

The next step is to understand how to customize your appliances' configuration. [Chapter 4, “Understanding the Email Pipeline”](#) provides a detailed overview of how email is routed through the system. Each feature is processed in order (from top to bottom) and is described in the remaining chapters of this guide.



Understanding the Email Pipeline

- [Overview of the Email Pipeline, page 4-1](#)
- [Email Pipeline Flows, page 4-1](#)
- [Incoming / Receiving, page 4-4](#)
- [Work Queue / Routing, page 4-7](#)
- [Delivery, page 4-10](#)

Overview of the Email Pipeline

The Email Pipeline is the flow of email as it is processed by the appliance. It has three phases:

- **Receipt** — As the appliance connects to a remote host to receive incoming email, it adheres to configured limits and other receipt policies. For example, verifying that the host can send your users mail, enforcing incoming connection and message limits, and validating the message's recipient.
- **Work Queue** — The appliance processes incoming and outgoing mail, performing tasks such as filtering, safelist/blocklist scanning, anti-spam and anti-virus scanning, Outbreak Filters, and quarantining.
- **Delivery** — As the appliance connects to send outgoing email, it adheres to configured delivery limits and policies. For example, enforcing outbound connection limits and processing undeliverable messages as specified.

Email Pipeline Flows

[Figure 4-1](#), [Figure 4-1](#), and [Figure 4-3](#) provide an overview of how email is processed through the system, from reception to routing to delivery. Each feature is processed in order (from top to bottom). You can test most of the configurations of features in this pipeline using the `trace` command.


Figure 4-1 Email Pipeline — Receiving Email Connections

Figure 4-2 Email Pipeline — Work Queue





Figure 4-3 Email Pipeline – Delivering Email

Incoming / Receiving

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

Related Topics

- [Host Access Table \(HAT\), Sender Groups, and Mail Flow Policies, page 4-5](#)
- [Received: Header, page 4-5](#)
- [Default Domain, page 4-5](#)
- [Bounce Verification, page 4-5](#)
- [Domain Map, page 4-6](#)
- [Recipient Access Table \(RAT\), page 4-6](#)
- [Alias Tables, page 4-6](#)

- [LDAP Recipient Acceptance, page 4-6](#)
- [SMTP Call-Ahead Recipient Validation, page 4-6](#)

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

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

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

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

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

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

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

For more information, see [Understanding Predefined Sender Groups and Mail Flow Policies, page 7-11](#).

Received: Header

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

For more information, see “Advanced Configuration Options” in the “Customizing Listeners” chapter.

Default Domain

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

For more information, see “SMTP Address Parsing Options” in the “Customizing Listeners” chapter.

Bounce Verification

Outgoing mail is tagged with a special key, and so if that mail is sent back as a bounce, the tag is recognized and the mail is delivered. For more information, see “Bounce Verification” in the “Configuring Routing and Delivery Features” chapter.

Domain Map

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

For more information, see “The Domain Map Feature” in the “Configuring Routing and Delivery Features” chapter.

Recipient Access Table (RAT)

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

For more information, see [Overview of Accepting or Rejecting Connections Based on the Recipient's Address, page 8-1](#).

Alias Tables

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

For more information about Alias Tables, see “Creating Alias Tables” in the “Configuring Routing and Delivery Features” chapter.

LDAP Recipient Acceptance

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

For more information, see the “LDAP Queries” chapter.

SMTP Call-Ahead Recipient Validation

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

For more information, see [Chapter 22, “Validating Recipients Using an SMTP Server.”](#)

Work Queue / Routing

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

**Note**

Data loss prevention (DLP) scanning is only available for outgoing messages. For information on where DLP message scanning occurs in the Work Queue, see [Message Splintering, page 10-5](#).

Related Topics

- [Email Pipeline and Security Services, page 4-7](#)
- [LDAP Recipient Acceptance, page 4-7](#)
- [Masquerading or LDAP Masquerading, page 4-8](#)
- [LDAP Routing, page 4-8](#)
- [Message Filters, page 4-8](#)
- [Email Security Manager \(Per-Recipient Scanning\), page 4-8](#)
- [Quarantines, page 4-10](#)

Email Pipeline and Security Services

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

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

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

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

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

LDAP Recipient Acceptance

You can use your existing LDAP infrastructure to define how the recipient email address of incoming messages (on a public listener) should be handled during the SMTP conversation or within the workqueue. See “Accept Queries” in the “Customizing Listeners” chapter. This allows the appliance to

combat directory harvest attacks (DHAP) in a unique way: the system accepts the message and performs the LDAP acceptance validation within the SMTP conversation or the work queue. If the recipient is not found in the LDAP directory, you can configure the system to perform a delayed bounce or drop the message entirely.

For more information, see the “LDAP Queries” chapter.

Masquerading or LDAP Masquerading

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

For more information about masquerading via a static mapping table, see “Configuring Masquerading” in the “Configuring Routing and Delivery Features” chapter.

For more information about masquerading via an LDAP query, see the “LDAP Queries” chapter.

LDAP Routing

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

For more information, see the “LDAP Queries” chapter.

Message Filters

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

For more information, see the “Using Message Filters to Enforce Email Policies” chapter.

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

Email Security Manager (Per-Recipient Scanning)

- [Safelist/Blocklist Scanning, page 4-9](#)
- [Anti-Spam, page 4-9](#)
- [Anti-Virus, page 4-9](#)
- [File Reputation Scanning and File Analysis, page 4-9](#)
- [Content Filters, page 4-10](#)
- [Outbreak Filters, page 4-10](#)

Safelist/Blocklist Scanning

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

Anti-Spam

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

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

See [Chapter 13, “Anti-Spam”](#) for more information.

Anti-Virus

Your appliance includes integrated virus scanning engines. You can configure the appliance to scan messages and attachments for viruses on a per-“mail policy” basis. You can configure the appliance to take actions such as the following when a virus is found:

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

Messages released from quarantines (see [Quarantines, page 4-10](#)) are scanned for viruses. See [Chapter 12, “Anti-Virus”](#) for more information about Anti-Virus scanning.

File Reputation Scanning and File Analysis

You can configure the appliance to scan message attachments for emerging and targeted threats. Available actions are similar to those for anti-virus scanning.

For more information, see [Chapter 16, “File Reputation Filtering and File Analysis.”](#)

Content Filters

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

For more information about Content Filters, see [Overview of Content Filters, page 11-1](#).

Outbreak Filters

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

Messages in the Outbreak quarantine are processed like any other message in a quarantine. For more information about quarantines and the Work Queue, see [Quarantines, page 4-10](#).

See [Chapter 14, “Outbreak Filters”](#) for more information.

Quarantines

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

The following Work Queue features can send messages to quarantines:

- Spam filters
- Message Filters
- Anti-Virus
- Outbreak Filters
- Content Filters
- File Analysis (Advanced Malware Protection)

Messages delivered from quarantines are re-scanned for threats.

Related Topics

- [Chapter 30, “Policy, Virus, and Outbreak Quarantines”](#)
- [Chapter 31, “Spam Quarantine”](#)

Delivery

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

Related Topics

- [Virtual gateways, page 4-11](#)

- [Delivery Limits, page 4-11](#)
- [Domain-Based Limits, page 4-11](#)
- [Domain-Based Routing, page 4-11](#)
- [Global Unsubscribe, page 4-11](#)
- [Bounce Limits, page 4-12](#)

Virtual gateways

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

For more information, see “Using Virtual Gateway™ Technology” in the “Configuring Routing and Delivery Features” chapter.

Delivery Limits

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

For more information, see “Set Email Delivery Parameters” in the “Configuring Routing and Delivery Features” chapter.

Domain-Based Limits

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

For more information, see “Controlling Email Delivery” in the “Configuring Routing and Delivery Features” chapter.

Domain-Based Routing

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

For more information, see “Routing Email for Local Domains” in the “Configuring Routing and Delivery Features” chapter.

Global Unsubscribe

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

For more information, see “Using Global Unsubscribe” in the “Configuring Routing and Delivery Features” chapter.

Bounce Limits

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

For more information about bounce profiles, see “Directing Bounced Email” in the “Configuring Routing and Delivery Features” chapter.



Configuring the Gateway to Receive Email

- [Overview of Configuring the Gateway to Receive Email, page 5-1](#)
- [Working with Listeners, page 5-2](#)
- [Configuring Global Settings for Listeners, page 5-6](#)
- [Listening for Connection Requests by Creating a Listener via the GUI, page 5-8](#)
- [Listening for Connection Requests by Creating a Listener via the CLI, page 5-13](#)
- [Enterprise Gateway Configuration, page 5-15](#)

Overview of Configuring the Gateway to Receive Email

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

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

You can create the following types of listeners:

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


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

- **Listener properties.** Define global properties that apply to all listeners, and properties specific to each listener. For example, you can specify the IP interface and port to use for a listener, and whether it is a public or private listener. For details on how to do this, see [Working with Listeners, page 5-2](#).

- **Which hosts that are allowed to connect to the listener.** Define a set of rules that control incoming connections from remote hosts. For example, you can define remote hosts and whether or not they can connect to the listener. For details on how to do this, see [Defining Which Hosts Are Allowed to Connect Using the Host Access Table \(HAT\)](#), page 7-1.
- **(Public listeners only) The local domains for which the listener accepts messages.** Define which recipients are accepted by the public listener. For example, if your organization uses the domain `currentcompany.com` and it previously used `oldcompany.com`, then you might accept messages for both `currentcompany.com` and `oldcompany.com`. For details on how to do this, see [Accepting or Rejecting Connections Based on Domain Name or Recipient Address](#), page 8-1.

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

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



Working with Listeners

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

You can define global settings that apply to all listeners. For more information, see [Configuring Global Settings for Listeners](#), page 5-6.

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


- You can define multiple listeners per configured IP interface, but each listener must use a different port.
- By default, listeners use SMTP as the mail protocol to service email connections. However, you can also configure the appliance to service email connections using Quick Mail Queuing Protocol (QMQP). Do this using the `listenerconfig` CLI command.
- Listeners support both Internet Protocol version 4 (IPv4) and version 6 (IPv6) addresses. You can use either protocol version or both on a single listener. The listener uses the same protocol version for mail delivery as the connecting host. For example, if the listener is configured for both IPv4 and IPv6 and connects to a host that uses IPv6, the listener uses IPv6. However, if the listener is configured to only use IPv6 addresses, it cannot connect to a host that is only using IPv4 addresses.
- At least one listener (with default values) is configured on the appliance after running the System Setup Wizard. However, when you create a listener manually, AsyncOS does not use these default SBRS values.
- **C170 appliances:** By default, the System Setup Wizard walks you through configuring one public listener for both receiving mail from the Internet and for relaying email from your internal network. That is, one listener can perform both functions.

- To help test and troubleshoot the appliance, you can create a “blackhole” type listener instead of a public or private listener. When you create a blackhole listener, you choose whether messages are written to disk or not before they are deleted. (See the “Testing and Troubleshooting” chapter for more information.) Writing messages to disk before deleting them can help you measure the rate of receiving and the speed of the queue. A listener that doesn’t write messages to disk can help you measure the pure rate of receiving from your message generation systems. This listener type is only available through the `listenerconfig` command in the CLI.

Figure 5-2 illustrates a typical email gateway configuration created by the System Setup Wizard on appliance models that have more than two Ethernet interfaces. Two listeners are created: a public listener to serve inbound connections on one interface and a private listener to serve outbound connections on a second IP interface.

Figure 5-3 illustrates a typical email gateway configuration created by the System Setup Wizard on appliance models that have only two Ethernet interfaces. One public listener on a single IP interface is created to serve both inbound and outbound connections.

Figure 5-2 *Public and Private Listeners on Appliance Models with More than Two Ethernet Interfaces*




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

Cisco Email Security appliance

IP interface PrivateNet sends messages to internal mail hosts.

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

Figure 5-3 Public Listener on Appliance Models with Only Two Ethernet Interfaces

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

Configuring Global Settings for Listeners

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

Procedure

- Step 1** Choose Network > Listeners.
- Step 2** Click **Edit Global Settings**.
- Step 3** Make changes to the settings defined in the following table.

Table 5-1 *Listener Global Settings*

Global Setting	Description
Maximum Concurrent Connections	Set the maximum number of concurrent connections for listeners. The default value is 300. If the listener accepts both IPv4 and IPv6 connections, the number of connections is divided between the two. For example, if the maximum concurrent connections is 300, then the sum of IPv4 and IPv6 connections cannot exceed 300.
Maximum Concurrent TLS Connections	Set the maximum concurrent TLS connections across all listeners combined. The default value is 100. If the listener accepts both IPv4 and IPv6 TLS connections, the number of connections is divided between the two. For example, if the maximum concurrent connections is 100, then the sum of IPv4 and IPv6 TLS connections cannot exceed 100.
Injection Counters Reset Period	Allows you to adjust when the injection control counters are reset. For very busy systems maintaining counters for a very large number of different IP addresses, configuring the counters to be reset more frequently (for example, every 15 minutes instead of every 60 minutes) will ensure that the data does not grow to an unmanageable size and impact system performance. The current default value is 1 hour. You can specify periods ranging from as little as 1 minute (60 seconds) to as long as 4 hours (14,400 seconds). See Injection Control Periodicity, page 7-25.
Timeout Period for Unsuccessful Inbound Connections	Set the length of time AsyncOS will allow an unsuccessful inbound connection to remain intact before closing it. An unsuccessful connection can be an SMTP conversation in which SMTP or ESMTP commands continue to be issued without a successful message injection occurring. When the specified timeout is reached, the behavior is to send an error and disconnect: “421 Timed out waiting for successful message injection, disconnecting.” A connection is considered unsuccessful until it successfully injects a message. Only available for SMTP connections on public listeners. The default value is 5 minutes.

Table 5-1 Listener Global Settings

Global Setting	Description
Total Time Limit for All Inbound Connections	Set the length of time AsyncOS will allow an inbound connection to remain intact before closing it. This setting is intended to preserve system resources by enforcing a maximum allowable connection time. Once about 80% of this maximum connection time is reached the following message is issued: “421 Exceeded allowable connection time, disconnecting.” The appliance will attempt to disconnect when the connection exceeds 80% of the maximum connection time in order to prevent disconnecting mid-message. It is likely that a problem is occurring with the inbound connection if it is open long enough to reach 80% of the maximum connection time. Keep this threshold in mind when specifying the time limit. Only available for SMTP connections on public listeners. The default value is 15 minutes.
Maximum size of subject	Messages having subject size within the specified limit will be accepted and any other messages will be rejected. If you set this value to 0, no limit is applied.
HAT delayed rejections	Configure whether to perform HAT rejection at the message recipient level. By default, HAT rejected connections will be closed with a banner message at the start of the SMTP conversation. When an email is rejected due to HAT “Reject” settings, AsyncOS can perform the rejection at the message recipient level (RCPT TO), rather than at the start of the SMTP conversation. Rejecting messages in this way delays the message rejection and bounces the message, allowing AsyncOS to retain more detailed information about the rejected messages. For example, you can see the mail from address and each recipient address of the message which is blocked. Delaying HAT rejections also makes it less likely that the sending MTA will perform multiple retries. When you enable HAT delayed rejection, the following behavior occurs: • The MAIL FROM command is accepted, but no message object is created. • All RCPT TO commands are rejected with text explaining that access to send e-mail is refused. • If the sending MTA authenticates with SMTP AUTH, they are granted a RELAY policy and are allowed to deliver mail as normal. Note Only configurable from the CLI <code>listenerconfig --> setup</code> command.

Step 4 Submit and commit your changes.

Related Topics

- [Settings for Messages Containing Multiple Encodings: localeconfig, page 5-8](#)

Settings for Messages Containing Multiple Encodings: localeconfig

You can set the behavior of AsyncOS regarding modifying the encoding of message headings and footers during message processing. This setting is not configured via the GUI. Instead, it is configured via the `localeconfig` in the CLI.

Listening for Connection Requests by Creating a Listener via the GUI

Procedure

- Step 1** Choose Network > Listener.
- Step 2** Click **Add Listener**.
- Step 3** Configure the settings defined in the following table.

Table 5-2 *Listener Settings*

Name	Unique nickname you supply for the listener, for future reference. The names you define for listeners are case-sensitive. AsyncOS will not allow you to create two identical listener names.
Type of Listener	Choose one of the following types of listeners: • Public. Public listeners contain default characteristics for receiving email from the Internet. • Private. Private listeners are intended to be used for private (internal) networks.
Interface	Choose a configured appliance IP interface and TCP port on which to create the listener. Depending on the version of the IP address used by the interface, the listener accepts connections from IPv4 addresses, IPv6 addresses or from both versions. By default, SMTP uses port 25 and QMQP uses port 628.
Bounce Profile	Select a bounce profile (bounce profiles created via the <code>bounceconfig</code> command in the CLI are available in the list, see Creating a New Bounce Profile, page 24-40).
Disclaimer Above	Select a disclaimer to attach above or below emails (disclaimers created via the Mail Policies > Text Resources page or the <code>textconfig</code> command in the CLI are available in the list, see the “Text Resources” chapter).
Disclaimer Below	Select a disclaimer to attach above or below emails (disclaimers created via the Mail Policies > Text Resources page or the <code>textconfig</code> command in the CLI are available in the list, see the “Text Resources” chapter).
SMTP Authentication Profile	Specify an SMTP Authentication profile.
Certificate	Specify a certificate for TLS connections to the listener (certificates added via the Network > Certificates page or the <code>certconfig</code> command in the CLI are available in the list, see Overview of Encrypting Communication with Other MTAs, page 23-1).

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

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

Setting	Description
Allow Partial Domains	If enabled, will allow partial domains. Partial domains can be no domain at all, or a domain with no dots. The following addresses are examples of partial domains: • foo • foo@ • foo@bar This option <i>must</i> be enabled in order for the Default Domain feature to work properly. Add Default Domain: A default domain to use for email addresses without a fully qualified domain name. This option is disabled unless Allow Partial Domains is enabled in SMTP Address Parsing options (see Listening for Connection Requests by Creating a Listener via the GUI, page 5-8). This affects how a listener modifies email that it relays by adding the “default sender domain” to sender and recipient addresses that do not contain fully-qualified domain names. (In other words, you can customize how a listener handles “bare” addresses). If you have a legacy system that sends email without adding (appending) your company’s domain to the sender address, use this to add the default sender domain. For example, a legacy system may automatically create email that only enters the string “joe” as the sender of the email. Changing the default sender domain would append “@yourdomain.com” to “joe” to create a fully-qualified sender name of joe@yourdomain.com.
Source Routing	Determines behavior if source routing is detected in the “MAIL FROM” and “RCPT TO” addresses. Source routing is a special form of an email address using multiple ‘@’ characters to specify routing (for example: @one.dom@two.dom:joe@three.dom). If set to “reject,” the address will be rejected. If “strip,” the source routing portion of the address will be deleted, and the message will be injected normally.
Unknown Address Literals	Determines behavior for when an address literal is received that the system cannot handle. Currently, this is everything except for IPv4. Thus, for example, for an IPv6 address literal, you can either reject it at the protocol level, or accept it and immediately hard bounce it. Recipient addresses containing literals will cause an immediate hard bounce. Sender addresses may get delivered. If the message cannot be delivered, then the hard bounce will hard bounce (double hard bounce). In the case of reject, both sender and recipient addresses will be rejected immediately at the protocol level.
Reject These Characters in User Names	Usernames that include characters (such as % or !, for example) entered here will be rejected.

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

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

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

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

For more information about creating LDAP queries, see [LDAP Queries, page 25-1](#).

Query Type	Description
Accept Queries	For Accept queries, select the query to use from the list. You can specify whether the LDAP Accept occurs during the work queue processing or during the SMTP conversation. For LDAP Accept during the work queue processing, specify the behavior for non-matching recipients: bounce or drop. For LDAP Accept during the SMTP conversation, specify how to handle mail if the LDAP server is unreachable. You can elect to allow messages or drop the connection with a code and custom response. Finally, select whether or not to drop connections if the Directory Harvest Attack Prevention (DHAP) threshold is reached during an SMTP conversation. Performing recipient validation in the SMTP conversation can potentially reduce the latency between multiple LDAP queries. Therefore, you might notice an increased load on your directory server when you enable conversational LDAP Accept. See Overview of LDAP Queries, page 25-1 for more information.
Routing Queries	For routing queries, select the query from the list. See Overview of LDAP Queries, page 25-1 for more information.
Masquerade Queries	For masquerade queries, select a query from the list, and select which address to masquerade, such as the From or CC header addresses. See Overview of LDAP Queries, page 25-1 for more information.
Group Queries	For group queries, select the query from the list. See Overview of LDAP Queries, page 25-1 for more information.

Step 7 Submit and commit your changes.

Related Topics

- [Partial Domains, Default Domains, and Malformed MAIL FROMs, page 5-12](#)

Partial Domains, Default Domains, and Malformed MAIL FROMs

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

These features are mutually exclusive.

Listening for Connection Requests by Creating a Listener via the CLI

Table 5-3 lists some of the `listenerconfig` subcommands used in the tasks involved in creating and editing listeners.

Table 5-3 *Tasks for Creating Listeners*

Tasks for Creating Listeners	Command(s) and Subcommands	Reference
Create a new listener	<code>listenerconfig -> new</code>	
Edit global settings for listeners	<code>listenerconfig -> setup</code>	Configuring Global Settings for Listeners, page 5-6
Specify a bounce profile for the listener	<code>bounceconfig, listenerconfig -> edit -> bounceconfig</code>	Creating a New Bounce Profile, page 24-40
Associate a disclaimer with the listener	<code>textconfig, listenerconfig -> edit -> setup -> footer</code>	
Configure an SMTP Authentication	<code>smtpauthconfig, listenerconfig -> smtpauth</code>	
Configure SMTP address parsing	<code>textconfig, listenerconfig -> edit -> setup -> address</code>	
Configure a default domain for the listener	<code>listenerconfig -> edit -> setup -> defaultdomain</code>	
Add a received header to email	<code>listenerconfig -> edit -> setup -> received</code>	
Change bare CR and LF characters to CRLF	<code>listenerconfig -> edit -> setup -> cleansmtp</code>	
Modify the Host Access Table	<code>listenerconfig -> edit -> hostaccess</code>	
Accept email for local domains or specific users (RAT) (public listeners only)	<code>listenerconfig -> edit -> rcptaccess</code>	
Encrypt conversations on listeners (TLS)	<code>certconfig, settls, listenerconfig -> edit</code>	Overview of Encrypting Communication with Other MTAs, page 23-1
Choose the certificate (TLS)	<code>listenerconfig -> edit -> certificate</code>	Overview of Encrypting Communication with Other MTAs, page 23-1

See [Chapter 24, “Configuring Routing and Delivery Features”](#) for information about email routing and delivery configurations.

Related Topics

[Advanced HAT Parameters, page 5-14](#)

Advanced HAT Parameters

Table 5-4 defines the syntax of advanced HAT parameters. Note that for the numeric values below, you can add a trailing **k** to denote kilobytes or a trailing **m** to denote megabytes. Values with no letters are considered bytes. Parameters marked with an asterisk support the variable syntax shown in Table 5-4.


Table 5-4 **Advanced HAT Parameter Syntax**

Parameter	Syntax	Values	Example Values
Maximum messages per connection	max_msgs_per_session	Number	1000
Maximum recipients per message	max_rcpts_per_msg	Number	10000 1k
Maximum message size	max_message_size	Number	1048576 20M
Maximum concurrent connections allowed to this listener	max_concurrency	Number	1000
SMTP Banner Code	smtp_banner_code	Number	220
SMTP Banner Text (*)	smtp_banner_text	String	Accepted
SMTP Reject Banner Code	smtp_banner_code	Number	550
SMTP Reject Banner Text (*)	smtp_banner_text	String	Rejected
Override SMTP Banner Hostname	use_override_hostname	on off default	default
	override_hostname	String	newhostname
Use TLS	tls	on off required	on
Use anti-spam scanning	spam_check	on off	off
Use virus scanning	virus_check	on off	off
Maximum Recipients per Hour	max_rcpts_per_hour	Number	5k
Maximum Recipients per Hour Error Code	max_rcpts_per_hour_code	Number	452
Maximum Recipients per Hour Text (*)	max_rcpts_per_hour_text	String	Too many recipients
Use SenderBase	use_sb	on off	on
Define SenderBase Reputation Score	sbrs[value1:value2]	-10.0- 10.0	sbrs[-10:-7.5]
Directory Harvest Attack Prevention: Maximum Invalid Recipients Per Hour	dhap_limit	Number	150

Enterprise Gateway Configuration

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

Figure 5-4 Public and Private Listeners for an Enterprise Gateway



In this configuration, at least two listeners are required:

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

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

Figure 5-4 shows one public listener (A) and one private listener (B) configured on the appliance in this Enterprise Gateway configuration.



Sender Reputation Filtering

- [Overview of Sender Reputation Filtering, page 6-1](#)
- [SenderBase Reputation Service, page 6-1](#)
- [Editing Sender Reputation Filtering Score Thresholds for a Listener, page 6-5](#)
- [Entering Low SBRS Scores in the Message Subject, page 6-7](#)

Overview of Sender Reputation Filtering

Sender reputation filtering is the first layer of spam protection, allowing you to control the messages that come through the email gateway based on senders' trustworthiness as determined by the Cisco SenderBase™ Reputation Service.

The appliance can accept messages from known or highly reputable senders — such as customers and partners — and deliver them directly to the end user without any content scanning. Messages from unknown or less reputable senders can be subjected to content scanning, such as anti-spam and anti-virus scanning, and you can also throttle the number of messages you are willing to accept from each sender. Email senders with the worst reputation can have their connections rejected or their messages bounced based on your preferences.



Note

File reputation filtering is a separate service. For information, see [Chapter 16, “File Reputation Filtering and File Analysis.”](#)

SenderBase Reputation Service

The Cisco SenderBase Reputation Service, using global data from the SenderBase Affiliate network, assigns a SenderBase Reputation Score to email senders based on complaint rates, message volume statistics, and data from public blacklists and open proxy lists. The SenderBase Reputation Score helps to differentiate legitimate senders from spam sources. You can determine the threshold for blocking messages from senders with low reputation scores.

The SenderBase Security Network website (www.senderbase.org) provides a global overview of the latest email and web-based threats, displays current email traffic volume by country, and allows you to look up reputation scores based on IP address, URI or Domain.

**Note**

The SenderBase Reputation Service is only available with a current anti-spam feature key.

Related Topics

- [SenderBase Reputation Score \(SBRS\), page 6-2](#)
- [How SenderBase Reputation Filters Work, page 6-3](#)
- [Recommended Settings for Different Sender Reputation Filtering Approaches, page 6-4](#)
- [Outbreak Filters, page 14-1](#)
- [Chapter 28, “Using Email Security Monitor”](#)

SenderBase Reputation Score (SBRS)

The SenderBase Reputation Score (SBRS) is a numeric value assigned to an IP address based on information from the SenderBase Reputation Service. The SenderBase Reputation Service aggregates data from over 25 public blacklists and open proxy lists, and combines this data with global data from SenderBase to assign a score from -10.0 to +10.0, as follows:


Score	Meaning
-10.0	Most likely to be a source of spam
0	Neutral, or not enough information to make a recommendation
+10.0	Most likely to be a trustworthy sender

The lower (more negative) the score, the more likely that a message is spam. A score of -10.0 means that this message is “guaranteed” to be spam, while a score of 10.0 means that the message is “guaranteed” to be legitimate.

Using the SBRS, you configure the appliance to apply mail flow policies to senders based on their trustworthiness. (You can also create message filters to specify “thresholds” for SenderBase Reputation Scores to further act upon messages processed by the system. For more information, refer to “[SenderBase Reputation Rule, page 9-35](#)” and “[Bypass Anti-Spam System Action, page 9-72](#).”)

Figure 6-1

The SenderBase Reputation Service




1. SenderBase affiliates send real-time, global data
2. Sending MTA opens connection with the appliance
3. 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. Cisco returns the response based on the SenderBase Reputation Score

How SenderBase Reputation Filters Work

Sender reputation filter technology aims to shunt as much mail as possible from the remaining security services processing that is available on the appliance. (See [Understanding the Email Pipeline, page 4-1.](#))

When sender reputation filtering is enabled, 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, sender reputation filters can reduce the load on the content filters by as much as 50%.

Figure 6-2 Sender Reputation Filtering Example

Recommended Settings for Different Sender Reputation Filtering Approaches

Depending on the objectives of your enterprise, you can implement a conservative, moderate, or aggressive approach.

Approach	Characteristics	Whitelist	Blacklist	Suspectlist	Unknownlist
Sender Base Reputation Score range:					
Conservative	Near zero false positives, better performance	7 to 10	-10 to -4	-4 to -2	-2 to 7
Moderate (Installation default)	Very few false positives, high performance	Sender Base Reputation Scores are not used.	-10 to -3	-3 to -1	-1 to +10
Aggressive	Some false positives, maximum performance. This option shunts the most mail away from Anti-Spam processing.	4 to 10	-10 to -2	-2 to -1	-1 to 4
Mail Flow Policy:					
All approaches		Trusted	Blocked	Throttled	Accepted

Editing Sender Reputation Filtering Score Thresholds for a Listener

Use this procedure if you want to change the default SenderBase Reputation Service (SBRS) score thresholds or add a sender group for reputation filtering.

**Note**

Other settings related to SBRS score thresholds, and Mail Flow Policy settings, are described in [Chapter 7, “Defining Which Hosts Are Allowed to Connect Using the Host Access Table \(HAT\).”](#)

Before You Begin

- If your appliance is set to receive mail from a local MX/MTA, identify upstream hosts that may mask the sender's IP address. See [Determining Sender IP Address In Deployments with Incoming Relays, page 13-15](#) for more information.
- Understand Sender Base Reputation Scores. See [Defining Sender Groups by SenderBase Reputation Score, page 7-6](#).
- Choose a filtering approach for your organization and note the recommended settings for that approach. See [Recommended Settings for Different Sender Reputation Filtering Approaches, page 6-4](#).

Procedure

-
- Step 1** Select **Mail Policies > HAT Overview**.
- Step 2** Select the public listener from the **Sender Groups (Listener)** menu.
- Step 3** Click the link for a sender group.
For example, click the “SUSPECTLIST” link.
- Step 4** Click **Edit Settings**.
- Step 5** Enter the range of SenderBase Reputation Scores for this sender group.
For example, for “WHITELIST,” enter the range 7.0 to 10.
- Step 6** Click **Submit**.
- Step 7** Repeat as needed for each sender group for this listener.
- Step 8** Commit changes.
-

Related Topics

- [Testing Sender Reputation Filtering Using the SBRS, page 6-6](#)
- [Monitoring the Status of the SenderBase Reputation Service, page 6-7](#)
- [Chapter 7, “Defining Which Hosts Are Allowed to Connect Using the Host Access Table \(HAT\)”](#)
- [How to Configure the Appliance to Scan Messages for Spam, page 13-2](#)

Testing Sender Reputation Filtering Using the SBRS

Unless you regularly receive a large portion of spam, or you have set up “dummy” accounts to specifically receive spam for your organization, it may be difficult to immediately test the SBRS policies you have implemented. However, if you add entries for reputation filtering with SenderBase Reputation Scores into a listener’s HAT as indicated in [Table 6-1](#), you will notice that a smaller percentage of inbound mail will be “unclassified.”

Test the policies using the `trace` command with an arbitrary SBRS. See [Debugging Mail Flow Using Test Messages: Trace, page 40-1](#). The `trace` command is available in the CLI as well as the GUI.

Table 6-1 Suggested Mail Flow Policies for Implementing the SBRS

Policy Name	Primary Behavior (Access Rule)	Parameters	Value
\$BLOCKED	REJECT	None	
\$THROTTLED	ACCEPT	Maximum messages / session: Maximum recipients / message: Maximum message size: Maximum concurrent connections: Use Spam Detection: Use TLS: Maximum recipients / hour: Use SenderBase:	10 20 1 MB 10 ON OFF 20 (recommended) ON
\$ACCEPTED (Public Listener)	ACCEPT	Maximum messages / session: Maximum recipients / message: Maximum message size: Maximum concurrent connections: Use Spam Detection: Use TLS: Use SenderBase:	1,000 1,000 100 MB 1,000 ON OFF ON
\$TRUSTED	ACCEPT	Maximum messages / session: Maximum recipients / message: Maximum message size: Maximum concurrent connections: Use Spam Detection: Use TLS: Maximum recipients / hour: Use SenderBase:	1,000 1,000 100 MB 1,000 OFF OFF -1 (disabled) OFF

**Note**

In the \$THROTTLED policy, the maximum recipients per hour from the remote host is set to 20 recipients per hour, by default. Note that this setting controls the maximum throttling available. You can increase the number of recipients to receive per hour if this parameter is too aggressive. For more information on Default Host Access policies, see [Understanding Predefined Sender Groups and Mail Flow Policies](#), page 7-11.

Monitoring the Status of the SenderBase Reputation Service

The SenderBase page in the Security Services menu displays the connection status and the timestamp of the most recent query from the appliance to the SenderBase Network Status Server and SenderBase Reputation Score Service. The SenderBase Reputation Score Service sends the SBRS scores to the appliance. The SenderBase Network Server sends the appliance information about the IP addresses, domains, and organizations that are sending mail to you. AsyncOS uses this data for its reporting and email monitoring features.

Figure 6-3 SenderBase Network Status on the SenderBase Page

SenderBase Network Status		
Type	Status	Last Status Check
SenderBase Network Server	up	Wed Sep 10 13:44:52 2008 PDT
SenderBase Reputation Score Service	up	Wed Sep 10 13:44:52 2008 PDT

The `sbstatus` command in CLI displays the same information.

Entering Low SBRS Scores in the Message Subject

Although Cisco recommends throttling, an alternate way to use the SenderBase Reputation Service is to modify the subject line of suspected spam messages. To do this, use the message filter shown in [Table 6-2](#). This filter uses the `reputation` filter rule and the `strip-header` and `insert-header` filter actions to replace the subject line of messages having a SenderBase Reputation Score lower than -2.0 with a subject line that includes the actual SenderBase Reputation Score represented as: `{Spam SBRS}`. Replace `listener_name` in this example with the name of your public listener. (The period on its own line is included so that you can cut and paste this text directly into the command line interface of the `filters` command.)

Table 6-2 Message Filter to Modify Subject Header with SBRS: Example 1

```
sbrs_filter:

if ((recv-inj == "listener_name" AND subject != "\\{Spam -?[0-9.]+\\}"))

{

    insert-header("X-SBRS", "$REPUTATION");

    if (reputation <= -2.0)

    {

        strip-header("Subject");
```

```
        insert-header("Subject", "$Subject \\{Spam $REPUTATION\\}");  
    }  
}  
.
```

Related Topic

- [Chapter 9, “Using Message Filters to Enforce Email Policies”](#).



Defining Which Hosts Are Allowed to Connect Using the Host Access Table (HAT)

- [Overview of Defining Which Hosts Are Allowed to Connect, page 7-1](#)
- [Defining Remote Hosts into Sender Groups, page 7-3](#)
- [Defining Access Rules for Email Senders Using Mail Flow Policies, page 7-8](#)
- [Understanding Predefined Sender Groups and Mail Flow Policies, page 7-11](#)
- [Handling Messages from a Group of Senders in the Same Manner, page 7-13](#)
- [Working with the Host Access Table Configuration, page 7-21](#)
- [Using a List of Sender Addresses for Incoming Connection Rules, page 7-22](#)
- [SenderBase Settings and Mail Flow Policies, page 7-23](#)
- [Verifying Senders, page 7-28](#)

Overview of Defining Which Hosts Are Allowed to Connect

For every configured listener, you must define a set of rules that control incoming connections from remote hosts. For example, you can define remote hosts and whether or not they can connect to the listener. AsyncOS allows you to define which hosts are allowed to connect to the listener using the Host Access Table (HAT).

The HAT maintains a set of rules that control incoming connections from remote hosts for a listener. Every configured listener has its own HAT. You configure HATs for both public and private listeners.

To control incoming connections from remote hosts, you define the following information:

- **Remote hosts.** Define the way in which a remote host attempts to connect to the listener. You group remote host definitions into *sender groups*. For example, you can define multiple remote hosts in a sender group by IP address and partial hostname. You can also define remote hosts by their SenderBase reputation score. For more information, see [Defining Remote Hosts into Sender Groups, page 7-3](#).
- **Access rules.** You can define whether the defined remote hosts in the sender group are allowed to connect to the listener and under what conditions. You define access rules using *mail flow policies*. For example, you can define that a particular sender group is allowed to connect to the listener, but only allow a maximum number of messages per connection. For more information, see [Defining Access Rules for Email Senders Using Mail Flow Policies, page 7-8](#)

Define which hosts are allowed to connect to the listener on the Mail Policies > HAT Overview page. [Figure 7-1](#) shows the HAT Overview with the sender groups and mail flow policies defined by default for a public listener.

Figure 7-1 Mail Policies > HAT Overview Page — Public Listener

HAT Overview

Find Senders

Find Senders that Contain this Text: Find

Sender Groups (Listener: IncomingMail (172.19.1.86:25))

Add Sender Group... Import HAT...

Order	Sender Group	SenderBase™ Reputation Score ?	Mail Flow Policy	Delete
1	WHITELIST		TRUSTED	
2	BLACKLIST		BLOCKED	
3	SUSPECTLIST		THROTTLED	
4	UNKNOWNLIST		ACCEPTED	
	ALL		ACCEPTED	

Edit Order... Export HAT...

Key: Custom Default

When a listener receives a TCP connection, it compares the source IP address against the configured sender groups. It evaluates the sender groups in the order listed on the HAT Overview page. When it finds a match, it applies the configured mail flow policy to the connection. If you have configured multiple conditions within a sender group, that sender group is matched if any of the conditions match.

When you create a listener, AsyncOS creates predefined sender groups and mail flow policies for the listener. You can edit the predefined sender groups and mail flow policies, and create new sender groups and mail flow policies. For more information, see [Understanding Predefined Sender Groups and Mail Flow Policies, page 7-11](#).

You can export all information stored in a Host Access Table to a file, and you can import Host Access Table information stored in a file into the appliance for a listener, overriding all configured Host Access Table information. For more information, see [Working with the Host Access Table Configuration, page 7-21](#).

Related Topics

- [Default HAT Entries, page 7-2](#)

Default HAT Entries

By default, the HAT is defined to take different actions depending on the listener type:

- **Public listeners.** The HAT is set to *accept* email from *all* hosts.
- **Private listeners.** The HAT is set up to *relay* email from the host(s) you specify, and reject all other hosts.

In the HAT Overview, the default entry is named “ALL.” You can edit the default entry by clicking the mail flow policy for the ALL sender group on the Mail Policies > HAT Overview page.

**Note**

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

Defining Remote Hosts into Sender Groups

You can define the way in which remote hosts attempt to connect to a listener. You group remote host definitions into sender groups. A sender group is a list of remote hosts defined for the purpose of handling email from those senders in the same way.

A sender group is a list of senders identified by:

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

For more information on the list of acceptable addresses in sender groups, see [Sender Group Syntax, page 7-4](#).

When an SMTP server attempts an SMTP connection with the appliance, the listener evaluates the sender groups in order and assigns the connection to a sender group when it matches *any* criterion in the sender group, such as SenderBase reputation score, domain, or IP address.

**Note**

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

Define sender groups on the Mail Policies > HAT Overview page.

Related Topics

- [Sender Group Syntax, page 7-4](#)
- [Sender Groups Defined by Network Owners, Domains, and IP Addresses, page 7-4](#)
- [Defining Sender Groups by SenderBase Reputation Score, page 7-6](#)
- [Sender Groups Defined by Querying DNS Lists, page 7-7](#)

Sender Group Syntax

Table 7-1 Defining Remote Hosts in the HAT: Sender Group Syntax

Syntax	Meaning
n:n:n:n:n:n:n	IPv6 address; does not need to include leading zeroes.
n:n:n:n:n:n:n-n:n:n:n:n:n:n:n n:n:n-n:n:n:n:n:n	Range of IPv6 addresses; does not need to include leading zeroes.
n.n.n.n	Full (complete) IPv4 Address
n.n.n. n.n.n n.n. n.n n. n	Partial IPv4 address
n.n.n.n-n n.n.n-n. n.n.n-n n.n-n. n.n-n n-n. n-n	Range of IPv4 addresses
yourhost.example.com	A fully-qualified domain name
.partialhost	Everything within the partialhost domain
n/c n.n/c n.n.n/c n.n.n.n/c	IPv4 CIDR address block
n:n:n:n:n:n:n/c	IPv6 CIDR address block; does not need to include leading zeroes
SBRS[n:n] SBRS[none]	SenderBase Reputation Score. For more information, see Defining Sender Groups by SenderBase Reputation Score, page 7-6 .
SBO:n	SenderBase Network Owner Identification Number. For more information, see Defining Sender Groups by SenderBase Reputation Score, page 7-6 .
dnslist[dnserver.domain]	DNS List query. For more information, see Sender Groups Defined by Querying DNS Lists, page 7-7 .
ALL	Special keyword that matches ALL addresses. This applies only to the ALL sender group, and is always included (but not listed).

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

Since the SMTP protocol has no built-in method for authenticating senders of email, senders of unsolicited bulk email have been successful at employing a number of tactics for hiding their identity. Examples include spoofing the Envelope Sender address on a message, using a forged HELO address,

or simply rotating through different domain names. This leaves many mail administrators asking themselves the fundamental question, “Who is sending me all of this email?” To answer this question, the SenderBase Reputation Service has developed a unique hierarchy for aggregating identity-based information based on the IP address of the connecting host — the one thing that is almost impossible for a sender to forge in a message.

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

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

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

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

Related Topics

- [Setting Policies Based on the HAT, page 7-5](#)

Setting Policies Based on the HAT

[Table 7-2](#) lists some examples of network owners and organizations.

Table 7-2 Example of Network Owners and Organizations

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

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

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

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

- **Which IP addresses are controlled by this sender?**

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

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

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

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

The Mail Flow Monitor feature uses its differentiation between SenderBase network owners and SenderBase organizations to determine how to allot connections per sender, based on logic in SenderBase. See the “Using Email Security Monitor” chapter for more information on using the Mail Flow Monitor feature.

Defining Sender Groups by SenderBase Reputation Score

The appliance can query the SenderBase Reputation Service to determine a sender’s reputation score (SBRs). The SBRs is a numeric value assigned to an IP address, domain, or organization based on information from the SenderBase Reputation Service. The scale of the score ranges from -10.0 to +10.0, as described in [Table 7-3](#).

Table 7-3 *Definition of the SenderBase Reputation Score*

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

Using the SBRS, you configure the appliance to apply mail flow policies to senders based on their trustworthiness. For example, all senders with a score less than -7.5 could be rejected. This is most easily accomplished via the GUI; see [Creating a Sender Group for Message Handling, page 7-13](#). However, if you are modifying an exported HAT in a text file, the syntax for including SenderBase Reputation Scores is described in [Table 7-4](#).

Table 7-4 Syntax for SenderBase Reputation Scores

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



Note

Network owners added to a HAT via the GUI use the syntax `SB0:n`, where *n* is the network owner's unique identification number in the SenderBase Reputation Service.

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



Note

You can also create message filters to specify “thresholds” for SenderBase Reputation Scores to further act upon messages processed by the system. For more information, see “SenderBase Reputation Rule,” “Bypass Anti-Spam System Action,” and “Bypass Anti-Virus System Action” in the anti-spam and anti-virus chapters.

Sender Groups Defined by Querying DNS Lists

You also have the ability in a listener's HAT to define a sender group as matching a query to a specific DNS List sever. The query is performed via DNS at the time of the remote client's connection. The ability to query a remote list also exists currently as a message filter rule (see “DNS List Rule” in the chapter on “Using Message Filters to Enforce Email Policies”), but only once the message content has been received in full.

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



Note

Some DNS Lists use variable responses (for example, “127.0.0.1” versus “127.0.0.2” versus “127.0.0.3”) to indicate various facts about the IP address being queried against. If you use the message filter DNS List rule (see “DNS List Rule” in the chapter on “Using Message Filters to Enforce Email Policies”), you can compare the result of the query against different values. However, specifying a DNS List server to be queried in the HAT only supports a Boolean operation for simplicity (that is, does the IP address appear in the list or not)

**Note**

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

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

Defining Access Rules for Email Senders Using Mail Flow Policies

Mail flow policies allow you to control or limit the flow of email messages from a sender to the listener during the SMTP conversation. You control SMTP conversations by defining the following types of parameters in the mail flow policy:

- Connection parameters, such as maximum number of messages per connection.
- Rate limiting parameters, such as maximum number of recipients per hour.
- Modify custom SMTP codes and responses communicated during the SMTP conversation.
- Enable spam detection.
- Enable virus protection.
- Encryption, such as using TLS to encrypt the SMTP connection.
- Authentication parameters, such as using DKIM to verify incoming mail.

Ultimately, mail flow policies perform one of the following actions on connections from remote hosts:

- **ACCEPT.** Connection is accepted, and email acceptance is then further restricted by listener settings, including the Recipient Access Table (for public listeners).
- **REJECT.** Connection is initially accepted, but the client attempting to connect gets a 4XX or 5XX SMTP status code. No email is accepted.

**Note**

You can also configure AsyncOS to perform this rejection at the message recipient level (RCPT TO), rather than at the start of the SMTP conversation. Rejecting messages in this way delays the message rejection and bounces the message, allowing AsyncOS to retain more detailed information about the rejected messages. This setting is configured from the CLI `listenerconfig > setup` command. For more information, see [Listening for Connection Requests by Creating a Listener via the CLI, page 5-13](#).

- **TCPREFUSE.** Connection is refused at the TCP level.
- **RELAY.** Connection is accepted. Receiving for any recipient is allowed and is not constrained by the Recipient Access Table.

- **CONTINUE.** The mapping in the HAT is ignored, and processing of the HAT continues. If the incoming connection matches a later entry that is not CONTINUE, that entry is used instead. The CONTINUE rule is used to facilitate the editing of the HAT in the GUI. For more information, see [Creating a Sender Group for Message Handling, page 7-13](#).

Related Topics

- [HAT Variable Syntax, page 7-9](#)

HAT Variable Syntax

[Table 7-5](#) defines a set of variables that can also be used in conjunction with the custom SMTP and Rate Limiting banners defined for a mail flow policy. Variable names are case-insensitive. (That is, `$group` is equivalent to `$Group`.)

Table 7-5 *HAT Variable Syntax*

Variable	Definition
<code>\$Group</code>	Replaced by the name of the sender group that was matched in the HAT. If the sender group has no name, “None” is displayed.
<code>\$Hostname</code>	Replaced by the remote hostname if and only if it has been validated by the 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.
<code>\$OrgID</code>	Replaced by the SenderBase Organization ID (an integer value). If the appliance cannot obtain a SenderBase Organization ID, or if the SenderBase Reputation Service did not return a value, “None” is displayed.
<code>\$RemoteIP</code>	Replaced by the IP address of the remote client.
<code>\$HATEntry</code>	Replaced by the entry in the HAT that the remote client matched.

Related Topics

- [Using HAT Variables, page 7-9](#)
- [Testing HAT Variables, page 7-10](#)

Using HAT Variables



Note

These variables can be used with the `smtp_banner_text` and `max_rcpts_per_hour_text` advanced HAT parameters described in the “Configuring the Gateway to Receive Email” chapter.

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

Figure 7-2 Using HAT Variables

Rate Limiting:	Max. Recipients Per Hour:	☒ Unlimited ☐
	Max. Recipients Per Hour Code:	452
	Max. Recipients Per Hour Text:	Too many recipients received this hour from Host: \$hostname

Or like this, in the CLI:

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

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

```
[220]> 200
```

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

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

Testing HAT Variables

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

```
# telnet IP_address_of_Email_Security_Appliance
```

```
220 hostname ESMTP
```

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

Understanding Predefined Sender Groups and Mail Flow Policies

Table 7-6 lists the predefined sender groups and mail flow policies that are configured when a public listener is created.

Table 7-6 *Predefined Sender Groups and Mail Flow Policies for Public Listeners*

Predefined Sender Group	Description	Default Configured Mail Flow Policy
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.	\$TRUSTED
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.	\$BLOCKED
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.	\$THROTTLED

Table 7-6 *Predefined Sender Groups and Mail Flow Policies for Public Listeners (continued)*

Predefined Sender Group	Description	Default Configured Mail Flow Policy
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 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. For more information on virus scanning engines, see Anti-Virus Scanning Overview, page 12-1 . For more information on the SenderBase Reputation Service, see SenderBase Reputation Service, page 6-1 .	\$ACCEPTED
ALL	Default sender group that applies to all other senders. For more information, see Default HAT Entries, page 7-2 .	\$ACCEPTED

[Table 7-7](#) lists the predefined sender groups and mail flow policies that are configured when a private listener is created.

Table 7-7 *Predefined Sender Groups and Mail Flow Policies for Private Listeners*

Predefined Sender Group	Description	Default Configured Mail Flow Policy
RELAYLIST	Add senders you know should be allowed to relay to the Relaylist sender group. The \$RELAYED mail flow policy is configured so that email from senders you are allowing to relay has no rate limiting, and the content from those senders is not scanned by the anti-spam scanning engine or anti-virus software. Note The RELAYLIST sender group includes the systems allowed to relay email when the System Setup Wizard was run.	\$RELAYED
ALL	Default sender group that applies to all other senders. For more information, see Default HAT Entries, page 7-2 .	\$BLOCKED

**Note**

When you run the System Setup Wizard on an appliance model that has only two Ethernet ports, you are prompted to create only one listener. It creates a public listener that also includes a \$RELAYED mail flow policy that is used to relay mail for internal systems. For appliance models that have more than two Ethernet ports, the RELAYLIST sender group and \$RELAYED mail flow policy only appear on private listeners.

Handling Messages from a Group of Senders in the Same Manner

Use the Mail Policies > HAT Overview and Mail Flow Policy pages to configure how the listener handles messages from senders. Do this by creating, editing, and deleting sender groups and mail flow policies.

Related Topics

- [Creating a Sender Group for Message Handling, page 7-13](#)
- [Adding a Sender to an Existing Sender Group, page 7-14](#)
- [Rearranging the Order of the Rules to Perform for Incoming Connections, page 7-14](#)
- [Searching for Senders, page 7-15](#)
- [Defining Rules for Incoming Messages Using a Mail Flow Policy, page 7-15](#)
- [Defining Default Values for Mail Flow Policies, page 7-20](#)

Creating a Sender Group for Message Handling

Procedure

-
- Step 1** Navigate to the **Mail Policies > HAT Overview** page.
- Step 2** Choose the listener to edit in the Listener field.
- Step 3** Click **Add Sender Group**.
- Step 4** Type the name of the sender group.
- Step 5** Select the order in which to place it in the list of sender groups.
- Step 6** (Optional) Enter a comment, for example information about this sender group or its settings.
- Step 7** Select a mail flow policy to which to apply this sender group.



Note If you do not know the mail flow policy you would like to apply to this group (or if no mail flow policies exist yet), then use the default “CONTINUE (no policy)” mail flow policy.

- Step 8** (Optional) Select a DNS list.
- Step 9** (Optional) Include senders for which SBRS has no information. This is referred to as “none” and generally denotes a suspect.
- Step 10** (Optional) Enter a DNS list.
- Step 11** (Optional) Configure host DNS verification settings.
For more information, see [Implementing Sender Verification — Example Settings, page 7-31](#).
- Step 12** Click **Submit and Add Senders** to create the group and begin adding senders to it.
- Step 13** Enter a sender using an IPv4 address, an IPv6 address, or a hostname. A sender can include a range of IP addresses and partial hostnames.



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

Step 14 (Optional) Enter a comment.

Step 15 Submit and commit your changes.

Related Topics

- [Editing Sender Reputation Filtering Score Thresholds for a Listener, page 6-5](#)

Adding a Sender to an Existing Sender Group

Procedure

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

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

Step 3 Submit and commit your changes.



Note When you add a domain to a sender group, two actual domains are listed in the GUI. For example, if you were adding the domain `example.net`, on the Add to Sender Group page, both `example.net` and `.example.net` are added. The second entry ensures that any host in the subdomain of `example.net` will be added to the sender group. For more information, see [Sender Group Syntax, page 7-4](#).



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

Step 4 Click **Save** to add the sender and return to the Incoming Mail Overview page.

Related Topics

- [Protecting Appliance-Generated Messages From the Spam Filter, page 13-14](#)
- [How to Configure the Appliance to Scan Messages for Spam, page 13-2](#)

Rearranging the Order of the Rules to Perform for Incoming Connections

If you add a sender group to a listener, you may need to edit the sender group order.

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.

Procedure

-
- Step 1** Navigate to the **Mail Policies > HAT Overview** page.
- Step 2** Choose the listener to edit in the Listener field.
- Step 3** Click **Edit Order**.
- Step 4** Type the new order for existing rows of sender groups in the HAT.
Cisco recommends maintaining the default order: RELAYLIST (certain hardware models only), followed by WHITELIST, BLACKLIST, SUSPECTLIST, and UNKNOWNLIST.
- Step 5** Submit and commit your changes.
-

Searching for Senders

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

Defining Rules for Incoming Messages Using a Mail Flow Policy

Consider the following rules and guidelines before creating a mail flow policy:

- Defaults for the policy are “greyed out” while the “Use Default” radio button is selected. To overwrite the default values, enable the feature or setting by selecting the “On” radio button and making changes to the now accessible values. To define default values, see [Defining Default Values for Mail Flow Policies, page 7-20](#).
- Some parameters depend on certain pre-configurations. (For example, the Directory Harvest Attack prevention setting requires that you have configured an LDAP Acceptance Query.)

Procedure

-
- Step 1** Navigate to the **Mail Policies > Mail Flow Policies** page.
- Step 2** Click **Add Policy**.
- Step 3** Enter the information described in [Table 7-8](#).

Table 7-8 *Mail Flow Policy Parameters*

Parameter	Description
Connections	
Maximum message size	The maximum size of a message that will be accepted by this listener. The smallest possible maximum message size is 1 kilobyte.
Maximum concurrent connections from a single IP	The maximum number of concurrent connections allowed to connect to this listener from a single IP address.
Maximum messages per connection	The maximum number of messages that can be sent through this listener per connection from a remote host.

Table 7-8 Mail Flow Policy Parameters (continued)

Parameter	Description
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. Note You can use some variables in this field. For more information, see HAT Variable Syntax, page 7-9 .
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 Limit for Hosts	
Max. Recipients per Hour	The maximum number of recipients per hour this listener will receive from a remote host. The number of recipients per sender IP address is tracked globally. Each listener tracks its own rate limiting threshold; however, because all listeners validate against a single counter, it is more likely that the rate limit will be exceeded if the same IP address (sender) is connecting to multiple listeners. Note You can use some variables in this field. For more information, see HAT Variable Syntax, page 7-9 .
Max. Recipients per Hour Code	The SMTP code returned when a host exceeds the maximum number of recipients per hour defined for this listener.
Max. Recipients Per Hour Exceeded Text	The SMTP banner text returned when a host exceeds the maximum number of recipients per hour defined for this listener.
Rate Limit for Sender	

Table 7-8 Mail Flow Policy Parameters (continued)

Parameter	Description
Max. Recipients per Time Interval	The maximum number of recipients during a specified time period that this listener will receive from a unique envelope sender, based on the mail-from address. The number of recipients is tracked globally. Each listener tracks its own rate limiting threshold; however, because all listeners validate against a single counter, it is more likely that the rate limit will be exceeded if messages from the same mail-from address are received by multiple listeners. Select whether to use the default maximum recipients, accept unlimited recipients, or specify another maximum number of recipients. Use the Default Mail Flow Policy settings to specify the maximum number of recipients and the time interval that will be used by the other mail flow policies by default. The time interval can only be specified using the Default Mail Flow Policy.
Sender Rate Limit Exceeded Error Code	The SMTP code returned when an envelope exceeds the maximum number of recipients for the time interval defined for this listener.
Sender Rate Limit Exceeded Error Text	The SMTP banner text returned when an envelope sender exceeds the maximum number of recipients for the time interval defined for this listener.
Exceptions	If you want certain envelope senders to be exempt from the defined rate limit, select an address list that contains the envelope senders. See Using a List of Sender Addresses for Incoming Connection Rules , page 7-22 for more information.
Flow Control	
Use SenderBase for Flow Control	Enable “look ups” to the SenderBase Reputation Service for this listener.
Group by Similarity of IP Addresses: (significant bits 0-32)	Used to track and rate limit incoming mail on a per-IP address basis while managing entries in a listener’s Host Access Table (HAT) in large CIDR blocks. You define a range of significant bits (from 0 to 32) by which to group similar IP addresses for the purposes of rate limiting, while still maintaining an individual counter for each IP address within that range. Requires “Use SenderBase” to be disabled. For more information about HAT significant bits, see “HAT Significant Bits Feature” in the “Configuring Routing and Delivery Features” chapter.
Directory Harvest Attack Prevention (DHAP)	
Directory Harvest Attack Prevention: Maximum Invalid Recipients Per Hour	The maximum number of invalid recipients per hour this listener will receive from a remote host. This threshold represents the total number of RAT rejections and SMTP call-ahead server rejections combined with the total number of messages to invalid LDAP recipients dropped in the SMTP conversation or bounced in the work queue (as configured in the LDAP accept settings on the associated listener). For more information on configuring DHAP for LDAP accept queries, see the “LDAP Queries” chapter.

Table 7-8 Mail Flow Policy Parameters (continued)

Parameter	Description
Directory Harvest Attack Prevention: Drop Connection if DHAP threshold is Reached within an SMTP Conversation	The 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	
TLS	Deny, Prefer, or Require Transport Layer Security (TLS) in SMTP conversations for this listener. If you select Preferred, you can make TLS mandatory for envelope senders from a specific domain or with a specific email address by selecting an Address List that specifies those domains and email addresses. When an envelope sender matching a domain or address in this list tries to send a message over a connection that does not use TLS, the appliance rejects the connection and the sender will have to try again using TLS. The Verify Client Certificate option directs the Email Security appliance to establish a TLS connection to the user’s mail application if the client certificate is valid. If you select this option for the TLS Preferred setting, the appliance still allows a non-TLS connection if the user doesn’t have a certificate, but rejects a connection if the user has an invalid certificate. For the TLS Required setting, selecting this option requires the user to have a valid certificate in order for the appliance to allow the connection. For information on creating an address list, see Using a List of Sender Addresses for Incoming Connection Rules, page 7-22. For information on using client certificates for TLS connections, see Establishing a TLS Connection from the Appliance, page 26-53.

Table 7-8 Mail Flow Policy Parameters (continued)

Parameter	Description
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.
If Both TLS and SMTP Authentication are enabled:	Require TLS to offer SMTP Authentication.
Domain Key Signing	
Domain Key/ DKIM Signing	Enable Domain Keys or DKIM signing on this listener (ACCEPT and RELAY only).
DKIM Verification	Enable DKIM verification.
S/MIME Decryption and Verification	
S/MIME Decryption/Verification	• Enable S/MIME decryption or verification. • Choose whether to retain or remove the digital signature from the messages after S/MIME verification. For triple wrapped messages, only the inner signature is retained or removed.
S/MIME Public Key Harvesting	
S/MIME Public Key Harvesting	Enable S/MIME public key harvesting.
Harvest Certificates on Verification Failure	Choose whether to harvest public keys if the verification of the incoming signed messages fail.
Store Updated Certificate	Choose whether to harvest updated public keys.
SPF/SIDF Verification	
Enable SPF/SIDF Verification	Enable SPF/SIDF signing on this listener. For more information, see the “Email Authentication” chapter.
Conformance Level	Set the SPF/SIDF conformance level. You can choose from SPF, SIDF or SIDF Compatible. For details, see the “Email Authentication” chapter.
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).
DMARC Verification	
Enable DMARC Verification	Enable DMARC verification on this listener. For more information, see DMARC Verification, page 20-35 .
Use DMARC Verification Profile	Select the DMARC verification profile that you want to use on this listener.

Table 7-8 Mail Flow Policy Parameters (continued)

Parameter	Description
DMARC Feedback Reports	Enable sending of DMARC aggregate feedback reports. For more information about DMARC aggregate feedback report, see DMARC Aggregate Reports, page 20-42. Note DMARC specification requires the feedback report messages to be DMARC compliant. Make sure that these messages are DKIM signed or you must publish appropriate SPF records.
Untagged Bounces	
Consider Untagged Bounces to be Valid	Applies only if bounce verification tagging (discussed in the “Configuring Routing and Delivery Features” chapter) 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 Verifying Senders, page 7-28 .
Exception Table	
Use Exception Table	Use the sender verification domain exception table. You can only have one exception table, but you can enable it per mail flow policy. See Sender Verification Exception Table, page 7-30 for more information.

**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 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.

Step 4 Submit and commit your changes.

Defining Default Values for Mail Flow Policies

Procedure

- Step 1** Click **Mail Policies > Mail Flow Policies**.
 - Step 2** Choose the listener to edit in the Listener field.
 - Step 3** Click the **Default Policy Parameters** link below the configured mail flow policies.
 - Step 4** Define the default values that all mail flow policies for this listener use.
- For more information on the properties, see [Defining Rules for Incoming Messages Using a Mail Flow Policy, page 7-15](#).

Step 5 Submit and commit your changes.

Working with the Host Access Table Configuration

You can export all information stored in a Host Access Table to a file, and you can import Host Access Table information stored in a file into the appliance for a listener, overwriting all existing Host Access Table information.

Related Topics

- [Exporting the Host Access Table Configuration to an External File, page 7-21](#)
- [Importing the Host Access Table Configuration from an External File, page 7-21](#)

Exporting the Host Access Table Configuration to an External File

Procedure

- Step 1** Navigate to the Mail Policies > HAT Overview page.
- Step 2** Choose the listener to edit in the Listener menu.
- Step 3** Click **Export HAT**.
- Step 4** Enter a file name for the exported HAT. This is the name of the file that will be created in the configuration directory on the appliance.
- Step 5** Submit and commit your changes.
-

Importing the Host Access Table Configuration from an External File

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

Procedure

- Step 1** Navigate to the Mail Policies > HAT Overview page.
- Step 2** Choose the listener to edit in the Listener menu.
- Step 3** Click **Import HAT**.
- Step 4** Select a file from the list.
-  **Note** The file to import must be in the configuration directory on the appliance.
-
- Step 5** Click **Submit**. You will see a warning message, asking you to confirm that you wish to remove all of the existing HAT entries.
- Step 6** Click **Import**.

Step 7 Commit your changes.

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

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

$BLOCKED

    REJECT {}

[ ... ]
```

Using a List of Sender Addresses for Incoming Connection Rules

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

Procedure

- Step 1** Select **Mail Policies > Address Lists**.
- Step 2** Click **Add Address List**.
- Step 3** Enter a name for the address list.
- Step 4** Enter a description of the address list.
- Step 5** (Optional) To enforce using full email addresses in the address list, select **Allow only full Email Addresses**.
- Step 6** Enter the addresses you want to include. You can use the following formats:

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



Note If you have selected **Allow only full Email Addresses**, you cannot use partial email addresses.

- IP address in their email address: `@[1.2.3.4]`
- All users in a domain: `@example.com`
- All users in a partial domain: `@.example.com`

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

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

Step 7 Submit and commit your changes.

SenderBase Settings and Mail Flow Policies

In order to classify connections to the appliance and apply mail flow policies (which may or may not contain rate limiting), a listener uses the following methodology:

Classification -> Sender Group -> Mail Flow Policy -> Rate Limiting

For more information, see [Sender Groups Defined by Network Owners, Domains, and IP Addresses, page 7-4](#).

The “Classification” stage uses the sending host’s IP address to classify an inbound SMTP session (received on a public listener) into a Sender Group. The Mail Flow Policy associated with that Sender Group may have parameters for rate limiting enabled. (Rate limiting limits the maximum number of messages per session, 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.)

Normally, in this process, recipients are counted against each sender in the corresponding named sender group. If mail is received from several senders in the same hour, the total recipients for all senders is compared against the limit.

There are some exceptions to this counting methodology:

- If the classification is done by Network Owner, then the SenderBase Reputation Service will automatically divide a large block of addresses into smaller blocks.
Counting of recipients and recipient rate limiting is done separately for each of these smaller blocks (usually, but not always, the equivalent of a /24 CIDR block).
- If the HAT Significant Bits feature is used. In this case, a large block of addresses may be divided into smaller blocks by applying the significant bits parameter associated with the policy.

Note that this parameter relates to the **Mail Flow Policy -> Rate Limiting** phase. It is not the same as the “bits” field in the “network/bits” CIDR notation that may be used to classify IP addresses in a Sender Group.

By default, SenderBase Reputation Service and IP Profiling support are *enabled* for public listeners and *disabled* for private listeners.

Related Topics

- [Timeouts for SenderBase Queries, page 7-23](#)
- [HAT Significant Bits Feature, page 7-24](#)

Timeouts for SenderBase Queries

When you configure a listener, you can determine how long the appliance caches information queried from the SenderBase Reputation Service. Then when you configure a mail flow policy, you can enable SenderBase to allow it to control the flow of mail into the listener.

Enable SenderBase in a mail flow policy in the GUI using the “Use SenderBase for Flow Control” setting when you configure a mail flow policy, or in the CLI using the `listenerconfig > hostaccess > edit` command.

HAT Significant Bits Feature

Beginning with the 3.8.3 release of AsyncOS, you can track and rate limit incoming mail on a per-IP address basis while managing sender group entries in a listener's Host Access Table (HAT) in large CIDR blocks. For example, if an incoming connection matched against the host "10.1.1.0/24," a counter could still be generated for each individual address within that range, rather than aggregating all traffic into one large counter.



Note

In order for the significant bits HAT policy option to take effect, you *must* not enable "User SenderBase" in the Flow Control options for the HAT (or, for the CLI, answer **no** to the question for enabling the SenderBase Information Service in the `listenerconfig -> setup` command: "Would you like to enable SenderBase Reputation Filters and IP Profiling support?"). That is, the Hat Significant Bits feature and enabling SenderBase IP Profiling support are mutually exclusive.

In most cases, you can use this feature to define sender groups *broadly* — that is, large groups of IP addresses such as "10.1.1.0/24" or "10.1.0.0/16" — while applying mail flow rate limiting *narrowly* to smaller groups of IP addresses.

The HAT Significant Bits feature corresponds to these components of the system:

- [HAT Configuration, page 7-24](#)
- [Significant Bits HAT Policy Option, page 7-24](#)
- [Injection Control Periodicity, page 7-25](#)

HAT Configuration

There are two parts of HAT configuration: sender groups and mail flow policies. 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. When using this feature, an IP address may be "classified in a CIDR block" (e.g. 10.1.1.0/24) sender group while being controlled as an individual host (/32). This is done via the "significant_bits" policy configuration setting.

Significant Bits HAT Policy Option

The HAT syntax allows for the `significant_bits` configuration option. When editing the default or a specific mail flow policy in a HAT (for example, when issuing the `listenerconfig -> edit -> hostaccess -> default` command) the following questions appear if:

- rate limiting is enabled, and
 - using SenderBase for flow control is disabled, or
 - Directory Harvest Attack Prevention (DHAP) is enabled for a mail flow policy (default or specific mail flow policy)

For example:

```
Do you want to enable rate limiting per host? [N]> y
```

```
Enter the maximum number of recipients per hour from a remote host.
```

[]> 2345

Would you like to specify a custom SMTP limit exceeded response? [Y]> n

Would you like to use SenderBase for flow control by default? [N]> n

Would you like to group hosts by the similarity of their IP addresses? [N]> y

Enter the number of bits of IP address to treat as significant, from 0 to 32.

[24]>

This feature also appears in the GUI in the Mail Policies > Mail Flow Policies page.

Figure 7-3 Enable the HAT Significant Bits Feature

Rate Limiting:	Max. Recipients Per Hour:	☒ Unlimited ☐
	Max. Recipients Per Hour Code:	452
	Max. Recipients Per Hour Text:	Too many recipients received this hour
Flow Control:	Use SenderBase for Flow Control:	☒ On ☐ Off
	Group by Similarity of IP Addresses:	<i>This Feature can only be used if Senderbase Flow Control is off.</i> ☒ Off ☐ (significant bits 0-32)

When the option to use SenderBase for flow control is set to “OFF” or Directory Harvest Attack Prevention is enabled, the “significant bits” value is applied to the connecting sender’s IP address, and the resulting CIDR notation is used as the token for matching defined sender groups within the HAT. Any rightmost bits that are covered by the CIDR block are “zeroed out” when constructing the string. Thus, if a connection from the IP address 1.2.3.4 is made and matches on a policy with the significant_bits option set to 24, the resultant CIDR block would be 1.2.3.0/24. So by using this feature, the HAT sender group entry (for example, 10.1.1.0/24) can have a different number of network significant bits (24) from the significant bits entry in the policy assigned to that group (32, in the example above).

Injection Control Periodicity

A global configuration option exists to allow you to adjust when the injection control counters are reset. For very busy systems maintaining counters for a very large number of different IP addresses, configuring the counters to be reset more frequently (for example, every 15 minutes instead of every 60 minutes) will ensure that the data does not grow to an unmanageable size and impact system performance.

The current default value is 3600 seconds (1 hour). You can specify periods ranging from as little as 1 minute (60 seconds) to as long as 4 hours (14,400 seconds).

Adjust this period via the GUI, using the global settings (for more information, see [Configuring Global Settings for Listeners, page 5-6](#)).

You can also adjust this period using the `listenerconfig -> setup` command in the CLI.

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

Currently configured listeners:

1. InboundMail (on PublicNet, 192.168.2.1) SMTP TCP Port 25 Public
2. OutboundMail (on PrivateNet, 192.168.1.1) SMTP TCP Port 25 Private

Choose the operation you want to perform:

- NEW - Create a new listener.
- EDIT - Modify a listener.
- DELETE - Remove a listener.
- SETUP - Change global settings.

```
[> setup
```

Enter the global limit for concurrent connections to be allowed across all listeners.

```
[300]>
```

Enter the global limit for concurrent TLS connections to be allowed across all listeners.

```
[100]>
```

Enter the maximum number of message header lines. 0 indicates no limit.

```
[1000]>
```

1. Allow SenderBase to determine cache time (Recommended)
2. Don't cache SenderBase data.
3. Specify your own cache time.

```
[1]> 3
```

Enter the time, in seconds, to cache SenderBase data:

```
[300]>
```

Enter the rate at which injection control counters are reset.

```
[1h]> 15m
```

Enter the timeout for unsuccessful inbound connections.

```
[5m]>
```

Enter the maximum connection time for inbound connections.

```
[15m]>
```

What hostname should Received: headers be stamped with?

1. The hostname of the Virtual Gateway(tm) used for delivering the message
2. The hostname of the interface the message is received on

```
[2]>
```

The system will always add a Message-ID header to outgoing messages that don't already have one. Would you like to do the same for incoming messages? (Not recommended.) [N]>

By default connections with a HAT REJECT policy will be closed with a banner message at the start of the SMTP conversation. Would you like to do the rejection at the message recipient level instead for more detailed logging of rejected mail? [N]>

Currently configured listeners:

1. InboundMail (on PublicNet, 192.168.2.1) SMTP TCP Port 25 Public
2. OutboundMail (on PrivateNet, 192.168.1.1) SMTP TCP Port 25 Private

```
[]>
```

Verifying Senders

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

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

The sender verification feature consists of the following components:

- **Verification of the connecting host.** This occurs prior to the SMTP conversation. For more information, see [Sender Verification: Host, page 7-28](#).
- **Verification of the domain portion of the envelope sender.** This occurs during the SMTP conversation. For more information, see [Sender Verification: Envelope Sender, page 7-29](#).

Related Topics

- [Sender Verification: Host, page 7-28](#)
- [Sender Verification: Envelope Sender, page 7-29](#)
- [Implementing Sender Verification — Example Settings, page 7-31](#)
- [Testing Your Settings for Messages from Unverified Senders, page 7-36](#)
- [Sender Verification and Logging, page 7-37](#)
- [Enabling Host DNS Verification via the CLI, page 7-38](#)

Sender Verification: Host

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

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

Unverified senders are classified into the following categories:

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

Using the sender group “Connecting Host DNS Verification” settings, you can specify a behavior for unverified senders (see [Throttling Messages from Unverified Senders Using the SUSPECTLIST Sender Group](#), page 7-32).

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

Sender Verification: Envelope Sender

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

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

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

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

For each mail flow policy, you can:

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

You can use the sender verification exception table to store a list of domains or addresses from which mail will be automatically allowed or rejected (see [Sender Verification Exception Table](#), page 7-30). The sender verification exception table can be enabled independently of Envelope Sender verification. So, for example, you can still reject special addresses or domains specified in the exception table without enabling envelope sender verification. You can also always allow mail from internal or test domains, even if they would not otherwise be verified.

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

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

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

Related Topics

- [Partial Domains, Default Domains, and Malformed MAIL FROMs, page 7-30](#)
- [Custom SMTP Code and Response, page 7-30](#)
- [Sender Verification Exception Table, page 7-30](#)

Partial Domains, Default Domains, and Malformed MAIL FROMs

If you enable envelope sender verification or disable allowing partial domains in SMTP Address Parsing options for a listener (see the SMTP Address Parsing Options section in the “Configuring the Gateway to Receive Email” chapter), the default domain settings for that listener will no longer be used.

These features are mutually exclusive.

Custom SMTP Code and Response

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

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

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

Sender Verification Exception Table

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

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

The sender verification exception table is defined in the GUI via the Mail Policies > Exception Table page (or the CLI, via the `exceptionconfig` command) and then is enabled on a per-policy basis via the GUI (see [Defining Messages to Send to Unverified Senders Using the ACCEPTED Mail Flow Policy, page 7-34](#)) or the CLI (see the *Cisco AsyncOS CLI Reference Guide*).

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

Figure 7-4 Exception Table Listing
Exception Table

Order	Exception	Behavior	SMTP Response	Delete
1	pres@whitehouse.gov	Allow	N/A	

See [Excluding Unverified Senders from Sender Verification Rules Based on Sender's Email Address, page 7-35](#) for more information about modifying the exception table.

Implementing Sender Verification — Example Settings

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

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

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

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

[Table 7-9](#) shows the suggested settings for implementing sender verification:

Table 7-9 Sender Verification: Suggested Settings

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

Related Topics

- [Throttling Messages from Unverified Senders Using the SUSPECTLIST Sender Group, page 7-32](#)
- [Implementing More Stringent Throttling Settings for Unverified Senders, page 7-33](#)
- [Defining Messages to Send to Unverified Senders Using the ACCEPTED Mail Flow Policy, page 7-34](#)
- [Excluding Unverified Senders from Sender Verification Rules Based on Sender's Email Address, page 7-35](#)
- [Searching for Addresses within the Sender Verification Exception Table, page 7-35](#)

Throttling Messages from Unverified Senders Using the SUSPECTLIST Sender Group

Procedure

- Step 1** Select **Mail Policies > HAT Overview**.
- Step 2** Click **SUSPECTLIST** in the list of sender groups.

Figure 7-5 *HAT Overview Page*

HAT Overview

The screenshot shows the 'Find Senders' section at the top with a search bar and a 'Find' button. Below is the 'Sender Groups (Listener: IncomingMail (172.19.0.86:25))' section. It includes an 'Add Sender Group...' button and an 'Import HAT...' button. The main table lists sender groups with their reputation scores and mail flow policies. The 'SUSPECTLIST' group is highlighted with a red box.

Order	Sender Group	SenderBase™ Reputation Score ?	Mail Flow Policy	Delete
1	WHITELIST	[Progress Bar]	TRUSTED	[Delete Icon]
2	BLACKLIST	[Progress Bar]	BLOCKED	[Delete Icon]
3	SUSPECTLIST	[Progress Bar]	THROTTLED	[Delete Icon]
4	UNKNOWNLIST	[Progress Bar]	ACCEPTED	[Delete Icon]
	ALL	[Progress Bar]	ACCEPTED	

Buttons at the bottom: 'Edit Order...', 'Export HAT...', 'Import HAT...'.

- Step 3** Click **Edit Settings**.

Figure 7-6 *Sender Group: SUSPECTLIST: Edit Settings*

The screenshot shows the 'Sender Group Settings' for the 'SUSPECTLIST' group. It includes a 'Comment' field, a 'Policy' dropdown set to 'THROTTLED', and an 'SBRS (Optional)' section with a range of -4.0 to -1.0 and a checkbox for 'Include SBRS Scores of "None"'. There is also a 'DNS Lists (Optional)' field and a 'Connecting Host DNS Verification' section with three checkboxes, the last of which is checked.

Buttons at the bottom: 'Cancel', 'Submit'.

- Step 4** Select the **THROTTLED** policy from the list.

- Step 5** Check the “Connecting host reverse DNS lookup (PTR) does not match the forward DNS lookup (A)” checkbox under Connecting Host DNS Verification.
- Step 6** Submit and commit your changes.
- Now, senders for which reverse DNS lookups fail will match the SUSPECTLIST sender group and will receive the default action from the THROTTLED mail flow policy.

**Note**

You can also configure host DNS verification via the CLI. See [Enabling Host DNS Verification via the CLI, page 7-38](#) for more information.

Implementing More Stringent Throttling Settings for Unverified Senders

Procedure

- Step 1** Create a new mail flow policy (for this example, it is named THROTTLEMORE) and configure it with more stringent throttling settings.
- On the Mail Flow Policies page, click **Add Policy**
 - Enter a name for the mail flow policy, and select Accept as the Connection Behavior.
 - Configure the policy to throttle mail.
 - Submit and commit your changes.
- Step 2** Create a new sender group (for this example, it is named UNVERIFIED) and configure it to use the THROTTLEMORE policy:
- On the HAT Overview page, click **Add Sender Group**

Figure 7-7 Add Sender Group: THROTTLEMORE

Add Sender Group to IncomingMail (192.168.0.1:25)

Sender Group Settings	
Name:	UNVERIFIED
Order:	5
Comment:	Throttle when host record is not in DNS
Policy:	THROTTLEMORE
SBRS (Optional):	☐ to ☐ ☐ Include SBRS Scores of "None" Recommended for suspected senders only.
DNS Lists (Optional):	?
Connecting Host DNS Verification:	☒ Connecting host PTR record does not exist in DNS. ☐ Connecting host PTR record lookup fails due to temporary DNS failure. ☐ Connecting host reverse DNS lookup (PTR) does not match the forward DNS lookup (A).

- Select the THROTTLEMORE policy from the list.
- Check the “Connecting host PTR record does not exist in DNS” checkbox under Connecting Host DNS Verification.

- d. Submit and commit your changes.

Figure 7-8 HAT Overview

HAT Overview

The screenshot shows the 'Find Senders' section with a search bar and a 'Find' button. Below it, the 'Sender Groups (Listener: IncomingMail (172.19.0.86:25))' section is visible. It includes an 'Add Sender Group...' button and an 'Import HAT...' button. A table lists five sender groups: WHITELIST, BLACKLIST, SUSPECTLIST, UNVERIFIED, and UNKNOWNLIST, each with a reputation score bar and a mail flow policy. The 'ALL' group is highlighted in yellow. At the bottom, there are 'Edit Order...' and 'Export HAT...' buttons, and a 'Key:' section with 'Custom' and 'Default' options.

Order	Sender Group	SenderBase™ Reputation Score	Mail Flow Policy	Delete
1	WHITELIST	[Score bar]	TRUSTED	[Delete icon]
2	BLACKLIST	[Score bar]	BLOCKED	[Delete icon]
3	SUSPECTLIST	[Score bar]	THROTTLED	[Delete icon]
4	UNVERIFIED	[Score bar]	THROTTLEMORE	[Delete icon]
5	UNKNOWNLIST	[Score bar]	ACCEPTED	[Delete icon]
	ALL	[Score bar]	ACCEPTED	

Defining Messages to Send to Unverified Senders Using the ACCEPTED Mail Flow Policy

Procedure

- Step 1** Select **Mail Policies > Mail Flow Policies**.
- Step 2** On the Mail Flow Policies page, click on the ACCEPTED mail flow policy.
- Step 3** Scroll to the bottom of the mail flow policy:

Figure 7-9 ACCEPTED Mail Flow Policy Envelope Sender DNS Verification Settings

The screenshot shows the 'Envelope Sender DNS Verification' settings. It includes three sections for defining SMTP codes and text for different types of envelope senders: Malformed Envelope Senders, Envelope Senders whose domain does not resolve, and Envelope Senders whose domain does not exist. Each section has a 'Use Default (Off)' button, an 'On' button, and an 'Off' button. The 'Use Exception Table' section at the bottom also has similar buttons.

Envelope Sender DNS Verification:	SMTP Code:	SMTP Text:
Malformed Envelope Senders:	553	#5.5.4 Domain required for sender address
Envelope Senders whose domain does not resolve:	451	#4.1.3 Domain of sender address <\$Envelo
Envelope Senders whose domain does not exist:	553	#5.1.8 Domain of sender address <\$Envelo

- Step 4** Select **On** to enable envelope sender DNS verification for this mail flow policy.
- Step 5** You may also define custom SMTP code and responses.
- Step 6** Enable the domain exception table by selecting **On** for “Use Domain Exception Table.”

Step 7 Submit and commit your changes.

Excluding Unverified Senders from Sender Verification Rules Based on Sender's Email Address

Procedure

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



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

Step 2 Click **Add Domain Exception** on the Mail Policies > Exception Table page.

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

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

Step 5 Submit and commit your changes.

Searching for Addresses within the Sender Verification Exception Table

Procedure

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

Step 2 Click **Find**.

Figure 7-10 Searching for Matching Entries in the Exception Table

Exception Table

The screenshot shows the 'Find Domain Exception' section with a search input containing 'mjones@partner.com' and a 'Find' button. Below is the 'Domain Exception Table' with an 'Add Domain Exception...' button and a table of results.

Order	Exception	Behavior	SMTP Response	Delete
1	pres@whitehouse.gov	Reject	553, Envelope sender <\$EnvelopeSender> rej...	
2	@partner.com	Allow	N/A	

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

Figure 7-11 Listing Matching Entries in the Exception Table**Exception Table**

Find Domain Exception

Search for Email Address:

Domain Exceptions Matching "mjones@partner.com"

Order	Exception	Behavior	SMTP Response	Delete
2	@partner.com	Allow	N/A	

Testing Your Settings for Messages from Unverified Senders

Now that you have configured sender verification settings, you can verify the behavior of your appliance. Note that testing DNS-related settings is beyond the scope of this document.

Related Topics

- [Sending a Test Message with a Malformed MAIL FROM Sender Address, page 7-36](#)
- [Sending a Message from an Address That is Excluded from Sender Verification Rules, page 7-37](#)

Sending a Test Message with a Malformed MAIL FROM Sender Address

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

Procedure

-
- Step 1** Open a Telnet session to your appliance.
- Step 2** Use SMTP commands to send a test message with a malformed MAIL FROM (something like “admin” without a domain).



Note If you have configured your appliance to use a default domain or to specifically allow partial domains when sending or receiving email or if you have enabled address parsing (see the “Configuring the Gateway to Receive Email” chapter) you may not be able to create, send, and receive an email with a missing or malformed domain.

- Step 3** Verify that the message is rejected.

```
# telnet IP_address_of_Email_Security_Appliance port

220 hostname ESMTPE
helo example.com

250 hostname
```



```
mail from: admin

553 #5.5.4 Domain required for sender address
```

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

Sending a Message from an Address That is Excluded from Sender Verification Rules

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

Procedure

- Step 1** Add the following address to the exception table with an “Allow” behavior: `admin@zzzaaazzz.com`
- Step 2** Commit your changes.
- Step 3** Open a Telnet session to your appliance.
- Step 4** Use SMTP commands to send a test message from the email address you entered in the sender verification exception table (`admin@zzzaaazzz.com`).
- Step 5** Verify that the message is accepted.

```
# telnet IP_address_of_Email_Security_Appliance port

220 hostname ESMTP

helo example.com

250 hostname

mail from: admin@zzzaaazzz.com

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

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

Sender Verification and Logging

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

Related Topics

- [Envelope Sender Verification, page 7-38](#)

Envelope Sender Verification

Malformed Envelope Senders:

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

Domain does not exist (NXDOMAIN):

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

Domain does not resolve (SERVFAIL):

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

Enabling Host DNS Verification via the CLI

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

Table 7-10 shows the types of unverified senders and the corresponding CLI setting:

Table 7-10 Sender Group Settings and Corresponding CLI Values

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



Accepting or Rejecting Connections Based on Domain Name or Recipient Address

- [Overview of Accepting or Rejecting Connections Based on the Recipient's Address, page 8-1](#)
- [Overview of the Recipient Access Table \(RAT\), page 8-2](#)
- [Accessing the RAT, page 8-2](#)
- [Editing the Default RAT Entry, page 8-2](#)
- [Domains and Users, page 8-3](#)

Overview of Accepting or Rejecting Connections Based on the Recipient's Address

AsyncOS uses a Recipient Access Table (RAT) for each public listener to manage accept and reject actions for recipient addresses. Recipient addresses include these:

- Domains
- Email addresses
- Groups of email addresses

The System Setup Wizard guides the administrator in configuring at least one public listener (with default values) on the appliance. Configuring a public listener during setup involves specifying default local domains or specific addresses to accept mail. These local domains or specific addresses are the first entries in the RAT for that public listener.

For each public listener, the default entry, “All Other Recipients”, rejects email from all recipients. The administrator defines all local domains for which the appliance accepts messages. Optionally, you can also define specific users for whom the appliance will accept or reject messages. AsyncOS allows you to define acceptable local domains and specific users using the Recipient Access Table (RAT).

You might need to configure a listener to accept messages for multiple domains. For example, if your organization uses the domain `currentcompanyname.com` and it previously used `oldcompanyname.com`, then you might accept messages for both `currentcompanyname.com` and `oldcompanyname.com`. In this case, include both local domains in the RAT for your public listener.

(Note: the Domain Map feature can map messages from one domain to another. See the Domain Map feature section of the “Configuring Routing and Delivery Features” chapter.)

Overview of the Recipient Access Table (RAT)

The Recipient Access Table defines which recipients are accepted by a public listener. At a minimum, the table specifies the address and whether to accept or reject it.

The Recipient Access Table (RAT) page shows a listing of the entries in the RAT including the order, default action, and whether or not the entry has been configured to bypass LDAP accept queries.

Accessing the RAT

GUI

Step 1 Navigate to **Mail Policies > Recipient Access Table (RAT)**.

CLI

Step 1 Use the `listenerconfig` command with the `edit > rcptaccess > new` subcommands.

Editing the Default RAT Entry

Before you begin

- Set up a public listener.
- Plan edits with caution, ensuring you do not create an *open relay* on the Internet. An open relay (sometimes called an “insecure relay” or a “third-party” relay) is an SMTP email server that allows third-party relay of email messages. By processing mail that is neither for — nor from — a local user, an open relay makes it possible for an unscrupulous sender to route large volumes of spam through your gateway. By default, the RAT *rejects* all recipients to prevent creation of an open relay.
- Note that you cannot delete the default entry from the RAT.

Procedure

Step 1 Navigate to **Mail Policies > Recipient Access Table (RAT)**.

Step 2 Click **All Other Recipients**.

Domains and Users

Modifying the Domains For Which to Accept Messages using the RAT

Use the Mail Policies > Recipient Access Table (RAT) page to configure the local domains and specific users for which the appliance accepts messages. On this page, you can perform the following tasks:

- Add, delete, and modify entries in the RAT.
- Change the order of the entries.
- Export RAT entries to a text file.
- Import RAT entries from a text file. Importing from a text file overwrites the existing entries.

Related Topics

- [Adding Domains and Users For Which to Accept Messages, page 8-3](#)
- [Rearranging the Order of Domains and Users in the Recipient Access Table, page 8-5](#)
- [Exporting the Recipient Access Table to an External File, page 8-5](#)
- [Importing the Recipient Access Table from an External File, page 8-6](#)

Adding Domains and Users For Which to Accept Messages

Procedure

-
Step 1	Navigate to the Mail Policies > Recipient Access Table (RAT) page.
Step 2	Choose the listener to edit in the Overview for Listener field.
Step 3	Click Add Recipient .
Step 4	Select an order for the entry.
Step 5	Enter the recipient address.
Step 6	Choose to accept or reject the recipient.
Step 7	(Optional) Choose to bypass LDAP acceptance queries for the recipient.
Step 8	(Optional) Use a custom SMTP response for this entry. a. Select Yes for Custom SMTP Response. b. Enter an SMTP response code and text. Include the SMTP response to the RCPT TO command for the recipient.
Step 9	(Optional) Choose to bypass throttling by selecting Yes for Bypass Receiving Control.
Step 10	Submit and commit your changes.
-

Related Topics

- [Defining Recipient Addresses, page 8-4](#)
- [Bypassing LDAP Accept for Special Recipients, page 8-4](#)
- [Bypassing Throttling for Special Recipients, page 8-5](#)

Defining Recipient Addresses

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

[IPv4 address]	Specific Internet Protocol version 4 (IPv4) address of the host. Note that the IP address must be between the “[]” characters.
[IPv6 address]	Specific Internet Protocol version 6 (IPv6) address of the host. Note that the IP address must be between the “[]” characters.
division.example.com	Fully-qualified domain name.
.partialhost	Everything within the “partialhost” domain.
user@domain	Complete email address.
user@	Anything with the given username.
user@[IP_address]	Username at a specific IPv4 or IPv6 address. Note that the IP address must be between the “[]” characters. Note that “user@IP_address” (without the bracket characters) is not a valid address. The system will append the brackets when it receives the message to create a valid address, which could affect whether a recipient is matched in the RAT.



Note

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

Bypassing LDAP Accept for Special Recipients

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

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

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

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

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

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

Bypassing Throttling for Special Recipients

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

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

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

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

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

Rearranging the Order of Domains and Users in the Recipient Access Table

Procedure

- Step 1** Navigate to the Mail Policies > Recipient Access Table (RAT) page.
- Step 2** Choose the listener to edit in the Overview for Listener field.
- Step 3** Click **Edit Order**.
- Step 4** Change the order by arranging the values in the Order column.
- Step 5** Submit and commit your changes.

Exporting the Recipient Access Table to an External File

Procedure

- Step 1** Navigate to the Mail Policies > Recipient Access Table (RAT) page.

- Step 2** Choose the listener to edit in the Overview for Listener field.
 - Step 3** Click **Export RAT**.
 - Step 4** Enter a file name for the exported entries.
This is the name of the file that will be created in the configuration directory on the appliance.
 - Step 5** Submit and commit your changes.
-

Importing the Recipient Access Table from an External File

When you import Recipient Access Table entries from a text file, all of the existing entries are removed from the Recipient Access Table.

Procedure

- Step 1** Navigate to the Mail Policies > Recipient Access Table (RAT) page.
- Step 2** Choose the listener to edit in the Overview for Listener field.
- Step 3** Click **Import RAT**.
- Step 4** Select a file from the list.
AsyncOS lists all text files in the configuration directory on the appliance.
- Step 5** Click **Submit**.
A warning message displays asking you to confirm that you want to remove all of the existing Recipient Access Table entries.
- Step 6** Click **Import**.
- Step 7** Commit your changes.
You can place “comments” in the file. Lines that begin with a ‘#’ character are considered comments and are ignored by AsyncOS. For example:

```
# File exported by the GUI at 20060530T220526

.example.com  ACCEPT

ALL  REJECT
```



Using Message Filters to Enforce Email Policies

The Cisco appliance contains extensive content scanning and message filtering technology that allows you to enforce corporate policies and act on specific messages as they enter or leave your corporate networks.

This chapter contains information about the powerful combinations of features available for policy enforcement: a content scanning engine, message filters, attachment filters, and content dictionaries.

This chapter contains the following sections:

- [Overview, page 9-1](#)
- [Components of a Message Filter, page 9-2](#)
- [Message Filter Processing, page 9-4](#)
- [Message Filter Rules, page 9-10](#)
- [Message Filter Actions, page 9-49](#)
- [Attachment Scanning, page 9-78](#)
- [Using the CLI to Manage Message Filters, page 9-89](#)
- [Message Filter Examples, page 9-107](#)
- [Configuring Scan Behavior, page 9-115](#)

Overview

Message filters allow you to create special rules describing how to handle messages as they are received by the Cisco appliance. A message filter specifies that a certain kind of email message should be given special treatment. Cisco message filters also allow you to enforce corporate email policy by scanning the content of messages for words you specify. This chapter contains the following sections:

- **Components of a message filter.** Message filters allow you to create special rules describing how to handle messages as they are received. Filter rules identify messages based on message or attachment content, information about the network, message envelope, message headers, or message body. Filter actions generate notifications or allow messages to be dropped, bounced, archived, blind carbon copied, or altered. For more information, see [Components of a Message Filter, page 9-2](#).
- **Processing Message Filters.** When AsyncOS processes message filters, the content that AsyncOS scans, the order of the processing, and the actions taken are based on several factors, including the message filter order, any prior processing that may have altered the message content, the MIME structure of the message, the threshold score configured for content matching, and structure of the query. For more information, see [Message Filter Processing, page 9-4](#).

- **Message Filter Rules.** Each filter has a rule that defines the collection of messages that the filter can act upon. You define those rules when you create a message filter. For more information, see [Message Filter Rules, page 9-10](#).
- **Message Filter Actions.** Each filter has an action that is performed on a message if the rule evaluates to `true`. There are two types of actions that can be performed: final actions (such as delivering, dropping, or bouncing a message), or non-final actions (such as stripping or inserting a header) which permit the message to be further processed. For more information, see [Message Filter Actions, page 9-49](#).
- **Attachment Scanning Message Filters.** Attachment scanning message filters allow you to strip attachments from messages that are inconsistent with your corporate policies, while still retaining the ability to deliver the original message. You can filter attachments based on their specific file type, fingerprint, or content. You can also scan image attachments using an image analyzer. The image analyzer creates algorithms to measure skin color, body size and curvature to determine the probability that the graphic contains inappropriate content. For more information, see [Attachment Scanning, page 9-78](#).
- **Using the CLI to Manage Message Filters.** The CLI accepts commands for working with message filters. For example, you might want to display, reorder, import or export a list of message filters. For more information, see [Using the CLI to Manage Message Filters, page 9-89](#).
- **Message Filter Examples.** This section contains some real world examples of filters with a brief discussion of each. For more information, see [Message Filter Examples, page 9-107](#).

Components of a Message Filter

Message filters allow you to create special rules describing how to handle messages as they are received. A message filter is comprised of message filter rules and message filter actions.

Related Topics

- [Message Filter Rules, page 9-2](#)
- [Message Filter Actions, page 9-2](#)
- [Message Filter Example Syntax, page 9-3](#)

Message Filter Rules

Message filter rules determine the messages that a filter will act on. Rules may be combined using the logical connectors AND, OR, and NOT to create more complex tests. Rule expressions may also be grouped using parentheses.

Message Filter Actions

The purpose of message filters is to perform actions on selected messages.

The two types of actions are:

- *Final* actions — such as `deliver`, `drop`, and `bounce` — end the processing of a message, and permit no further processing through subsequent filters.
- *Non-final* actions perform an action which permits the message to be processed further.

**Note**

Non-final message filter actions are cumulative. If a message matches multiple filters where each filter specifies a different action, then all actions are accumulated and enforced. However, if a message matches multiple filters specifying the same action, the prior actions are overridden and the final filter action is enforced.

Message Filter Example Syntax

The intuitive meaning of a filter specification is:

if the message matches the rule, *then* apply the actions in sequence. If the `else` clause is present, the actions within the `else` clause are executed in the event the message does *not match* the rule.

The name of the filter you specify makes it easier to manage filters when you are activating, deactivating, or deleting them.

Message filters use the following syntax:

Example Syntax	Purpose
<code>expedite:</code>	filter name
<code>if (recv-listener == 'InboundMail' or recv-int == 'notmain')</code>	rule specification
<code>{</code> <code>alt-src-host('outbound1');</code> <code>skip-filters();</code> <code>}</code>	action specification
<code>else</code> <code>{</code> <code>alt-src-host('outbound2');</code> <code>}</code>	optional alternative action specification

Note that you can omit any alternative actions:

Example Syntax	Purpose
<code>expedite2:</code>	filter name
<code>if ((not (recv-listener == 'InboundMail')) and (not (recv-int == 'notmain')))</code>	rule specification
<code>{</code> <code>alt-src-host('outbound2');</code> <code>skip-filters();</code> <code>}</code>	action specification

You can combine several filters in sequence within a single text file, one following the other.

You must enclose the values in filters in either single or double quotation marks. Single or double quotation marks must be equally paired on each side of the value; for example, the expressions `notify('customercare@example.com')` and `notify("customercare@example.com")` are both legal, but the expression `notify(customercare@example.com)` causes a syntax error.

Lines beginning with a '#' character are considered comments and are ignored; however, they are not preserved by AsyncOS as can be verified by viewing a filter via `filters -> detail`.

Message Filter Processing

When AsyncOS processes message filters, the content that AsyncOS scans, the order of the processing, and the actions taken are based on several factors:

- **Message filter order.** Message filters are maintained in an ordered list. When a message is processed, AsyncOS applies each message filter in the order it appears in the list. If a final action occurs, no further action is taken on the message. For more information, see [Message Filter Order, page 9-4](#).
- **Prior processing.** Actions performed on AsyncOS messages may add or remove headers before the message filter is evaluated. AsyncOS processes the message filter process on the headers that are present in the message at the time of processing. For more information, see [Message Header Rules and Evaluation, page 9-5](#).
- **The MIME structure of the message.** The MIME structure of the message determines which part of the message is treated as “body,” and which part of the message is treated as an “attachment”. Many message filters are configured to act on just the body or just the attachment part of the message. For more information, see [Message Bodies vs. Message Attachments, page 9-5](#).
- **The threshold score configured for the regular expression.** When you match a regular expression, you configure a “score” to tally up the number of times a match must occur before a filter action is taken. This allows you to “weight” the responses to different terms. For more information, see [Thresholds for Matches in Content Scanning, page 9-6](#).
- **The structure of the query.** When evaluating AND or OR tests within message filters, AsyncOS does not evaluate unneeded tests. In addition, it is important to note that the system does not evaluate the tests from left to right. Instead, when AND and OR tests are evaluated, the least expensive test is evaluated first. For more information, see [AND Test and OR Tests in Message Filters, page 9-9](#).

Related Topics

- [Message Filter Order, page 9-4](#)
- [Message Header Rules and Evaluation, page 9-5](#)
- [Message Bodies vs. Message Attachments, page 9-5](#)
- [Thresholds for Matches in Content Scanning, page 9-6](#)
- [AND Test and OR Tests in Message Filters, page 9-9](#)

Message Filter Order

Message filters are kept in an ordered list and numbered by their position in the list. When a message is processed, the message filters are applied in the associated numeric order. Therefore, filter number 30 will not have a chance to alter the source host of a message if filter number 9 has already executed a final action on (for example, bounced) the message. The position of a filter in the list can be changed via the system user interfaces. Filters imported via a file are ordered based on their relative order in the imported file.

After a final action, no further actions may be taken on the message.

Although a message may match a filter rule, the filter may not act upon that message for any of the following reasons:

- The filter is inactive.
- The filter is invalid.

- The filter has been superseded by an earlier filter that executed a final action for the message.

Message Header Rules and Evaluation

Filters evaluate “processed” headers rather than the original message headers when applying header rules. Thus:

- If a header was added by a previous processing action, it can now be matched by any subsequent header rule.
- If a header was stripped by a previous processing action, it can no longer be matched by any subsequent header rule.
- If a header was modified by a previous processing action, any subsequent header rule will evaluate the modified header and not the original message header.


This behavior is common to both message filters and content filters.

Message Bodies vs. Message Attachments

An email message is composed of multiple parts. Although RFCs define everything that comes after a message’s headers as a multipart “message body,” many users still conceptualize a message’s “body” and its “attachment” differently. When you use any of the Cisco message filters named *body-variable* or *attachment-variable*, the Cisco appliance attempts to distinguish the parts that most users consider to be the “body” and the “attachment” in the same way that many MUAs attempt to render these parts differently.

For the purposes of writing *body-variable* or *attachment-variable* message filter rules, everything after the message headers is considered the message *body*, whose content is considered the first text part of the MIME parts that are within the body. Anything after the content, (that is, any additional MIME parts) is considered an *attachment*. AsyncOS evaluates the different MIME parts of the message, and identifies the parts of the file that is treated as an attachment.

For example, [Figure 9-1](#) shows a message in the Microsoft Outlook MUA where the words “Document attached below.” appear as a plain text message *body* and the document “This is a Microsoft Word document.doc” appears as an *attachment*. Because many users conceptualize email this way (rather than as a multipart message whose first part is plain text and whose second part is a binary file), the Cisco uses the term “attachment” in message filters to create rules to differentiate and act on the .doc file part (in essence, the second MIME part) as opposed to the “body” of the message (the first, plain text part) — although, according to the language used in RFCs 1521 and 1522, a message’s *body* is comprised of all MIME parts.

Figure 9-1 Message with “Attachment”

Because the Cisco appliance makes this distinction between the *body* and the *attachment* in multipart messages, there are several cases you should be aware of when using the `body-variable` or `attachment-variable` message filter rules in order to achieve the expected behavior:

- If you have a message with a single text part—that is, a message containing a header of “Content-Type: text/plain” or “Content-Type: text/html” — the Cisco appliance will consider the entire message as the body. If the content type is anything different, the Cisco appliance considers it to be a single attachment.
- Some encoded files (uuencoded, for example) are included in the body of the email message. When this occurs, the encoded file is treated as an attachment, and it is extracted and scanned, while the remaining text is considered to be the body of the text.
- A single, non-text part is always considered an *attachment*. For example, a message consisting of only a .zip file is considered an attachment.

Thresholds for Matches in Content Scanning

When you add filter rules that search for patterns in the message body or attachments, you can specify the minimum threshold for the number of times the pattern must be found. When AsyncOS scans the message, it totals the “score” for the number of matches it finds in the message and attachments. If the minimum threshold is not met, the regular expression does not evaluate to true. You can specify this threshold for the following filter rules:

- body-contains
- only-body-contains
- attachment-contains
- every-attachment-contains
- dictionary-match
- attachment-dictionary-match

You can also specify a threshold value for the `drop-attachments-where-contains` action.



Note

You cannot specify thresholds for filter rules that scan headers or envelope recipients and senders.

Related Topics

- [Threshold Syntax, page 9-7](#)
- [Threshold Scoring for Message Bodies and Attachments, page 9-7](#)
- [Threshold Scoring Multipart/Alternative MIME Parts, page 9-7](#)

- [Threshold Scoring for Content Dictionaries, page 9-8](#)

Threshold Syntax

To specify a threshold for the minimum number of occurrences, specify the pattern and the minimum number of matches required to evaluate to true:

```
if(<filter rule>(<pattern>,<minimum threshold>)){
```

For example, to specify that the `body-contains` filter rule must find the value “Company Confidential” at least two times, use the following syntax:

```
if(body-contains('Company Confidential',2)){
```

By default, when AsyncOS saves a content scanning filter, it compiles the filter and assigns a threshold value of 1, if you have not assigned a value.

You can also specify a minimum number of pattern matches for values in a content dictionary. For more information about content dictionaries, see the “Text Resources” chapter.

Threshold Scoring for Message Bodies and Attachments

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

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

Threshold Scoring Multipart/Alternative MIME Parts

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

For example, you configure a `body-contains` filter rule and set the threshold to 4. You then receive a message that contains both plain text, HTML and two attachments. The message would use the following structure:

```
multipart/mixed
```

```
    multipart/alternative
```

```

text/plain

text/html

application/octet-stream

application/octet-stream

```

The `body-contains` filter rule would determine the score for this message by first scoring the `text/plain` and `text/html` parts of the message. It would then compare the results of these scores and select the highest score from the results. Next, it would add this result to the score from each of the attachments to determine the final score. Suppose the message has the following number of matches:

```

multipart/mixed

multipart/alternative

    text/plain (2 matches)

    text/html (2 matches)

application/octet-stream (1 match)

application/octet-stream

```

Because AsyncOS compares the matches for the `text/plain` and `text/html` parts, it returns a score of 3, which does not meet the minimum threshold to trigger the filter rule.

Threshold Scoring for Content Dictionaries

When you use a content dictionary, you can “weight” terms so that certain terms trigger filter actions more easily. For example, you may want not want to trigger a message filter for the term, “bank.” However, if the term, “bank” is combined with the term, “account,” and accompanied with an ABA routing number, you may want to trigger a filter action. To accomplish this, you can use a weighted dictionary to give increased importance to certain terms or a combination of terms. When a message filter that uses a content dictionary scores the matches for filter rule, it uses these weights to determine the final score. For example, suppose you create a content dictionary with the following contents and weights:

Table 9-1 **Sample Content Dictionary**

Term/Smart Identifier	Weight
ABA Routing Number	3
Account	2
Bank	1

When you associate this content dictionary with a `dictionary-match` or `attachment-dictionary-match` message filter rule, AsyncOS would add the weight for the term to the total “score” for each instance of the matching term found in the message. For example, if the message contains three instances of the term, “account” in the message body, AsyncOS would add a value of 6 to

the total score. If you set the threshold value for the message filter to 6, AsyncOS would determine that the threshold score has been met. Or, if the message contained one instance of each term, the total value would be 6, and this score would trigger the filter action.

AND Test and OR Tests in Message Filters

When evaluating AND or OR tests within message filters, AsyncOS does not evaluate unneeded tests. So, for example, if one side of an AND test is false, the system will not evaluate the other side. It is important to note that the system does not evaluate the tests from left to right. Instead, when AND and OR tests are evaluated, the least expensive test is evaluated first. For example, in the following filter, the `remote-ip` test will always be processed first because it has a lower cost than the `rcpt-to-group` test (generally LDAP tests are more expensive):

```
andTestFilter:

if (remote-ip == "192.168.100.100" AND rcpt-to-group == "GROUP")

    { ... }
```

Because the least expensive test is performed first, switching the order of the items in the test will have no effect. If you want to guarantee the order in which tests are performed, use nested `if` statements. This is also the best way to ensure that an expensive test is avoided whenever possible:

```
expensiveAvoid:

if (<simple tests>)

    { if (<expensive test>)

        { <action> }

    }
```

In a somewhat more complicated example, consider:

```
if (test1 AND test2 AND test3) { ... }
```

The system groups the expression from left to right, so this becomes:

```
if ((test1 AND test2) AND test3) { ... }
```

This means the first thing the system does is compare the cost of `(test1 AND test2)` against the cost of `test3`, evaluating the second AND first. If all three tests have the same cost, then `test3` will be performed first because `(test1 AND test2)` would be twice as expensive.

Message Filter Rules

Each message filter contains a rule that defines the collection of messages that a filter can act upon. You define the filter rules, and then you define a filter action for messages that return `true`.

Related Topics

- [Filter Rules Summary Table, page 9-10](#)
- [Regular Expressions in Rules, page 9-17](#)
- [Smart Identifiers, page 9-21](#)
- [Description and Examples of Message Filter Actions, page 9-58](#)

Filter Rules Summary Table

[Table 9-2](#) summarizes the rules you can use in message filters.

Table 9-2 *Message Filter Rules*

Rule	Syntax	Description
Subject Header	<code>subject</code>	Does the subject header match a certain pattern? See Subject Rule, page 9-24 .
Body Size	<code>body-size</code>	Is the body size within some range? See Body Size Rule, page 9-27 .
Envelope Sender	<code>mail-from</code>	Does the Envelope Sender (i.e., the Envelope From, <MAIL FROM>) match a given pattern? See Envelope Sender Rule, page 9-25 .
Envelope Sender in Group	<code>mail-from-group</code>	Is the Envelope Sender (i.e., the Envelope From, <MAIL FROM>) in a given LDAP group? See Envelope Sender in Group Rule, page 9-26 .
Sender Group	<code>sendergroup</code>	Which sender group was matched in a listener's Host Access Table (HAT)? See Sender Group Rule, page 9-26 .
Envelope Recipient	<code>rcpt-to</code>	Does the Envelope Recipient, (i.e. the Envelope To, <RCPT TO>) match a given pattern? See Envelope Recipient Rule, page 9-24 . Note: The <code>rcpt-to</code> rule is message-based. If a message has multiple recipients, only one recipient has to match the rule for the specified action to affect the message to all recipients.

Table 9-2 **Message Filter Rules**

Rule	Syntax	Description
Envelope Recipient in Group	<code>rcpt-to-group</code>	Is the Envelope Recipient, (i.e. the Envelope To, <RCPT TO>) in a given LDAP group? See Envelope Recipient in Group Rule, page 9-25 . Note: The <code>rcpt-to-group</code> 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.
Remote IP	<code>remote-ip</code>	Was the message sent from a remote host that matches a given IP address or IP block? See Remote IP Rule, page 9-27 .
Receiving Interface	<code>recv-int</code>	Did the message arrive via the named receiving interface? See Receiving IP Interface Rule, page 9-28 .
Receiving Listener	<code>recv-listener</code>	Did the message arrive via the named listener? See Receiving Listener Rule, page 9-28 .
Date	<code>date</code>	Is current time before or after a specific time and date? See Date Rule, page 9-28 .
Header	<code>header(<string>)</code>	Does the message contain a specific header? Does the value of that header match a certain pattern? See Header Rule, page 9-29 .
Random	<code>random(<integer>)</code>	Is a random number in some range? See Random Rule, page 9-30 .
Recipient Count	<code>rcpt-count</code>	How many recipients is this email going to? See Recipient Count Rule, page 9-30 .
Address Count	<code>addr-count()</code>	What is the cumulative number of recipients? This filter differs from the <code>rcpt-count</code> filter rule in that it operates on the message body headers instead of the envelope recipients. See Address Count Rule, page 9-31 .
SPF Status	<code>spf-status</code>	What was the SPF verification status? This filter rule allows you to query for different SPF verification results. You can enter a different action for each valid SPF/SIDF return value. See SPF-Status Rule, page 9-37 .
SPF Passed	<code>spf-passed</code>	Did the SPF/SIDF verification pass? This filter rule generalizes the SPF/SIDF results as a Boolean value. See SPF-Passed Rule, page 9-39 .
S/MIME Gateway Message	<code>smime-gateway</code>	Is the message S/MIME signed, encrypted, or signed and encrypted? See S/MIME Gateway Message Rule, page 9-39 .

Table 9-2 Message Filter Rules

Rule	Syntax	Description
S/MIME Gateway Verified	smime-gateway-verified	Is the S/MIME message successfully verified, decrypted, or decrypted and verified? See S/MIME Gateway Verified Rule, page 9-40
Image verdict	image-verdict	What was the image scanning verdict? This filter rule allows you to query for different image analysis verdicts. See Image Analysis, page 9-80 .
Workqueue count	workqueue-count	Is the work queue count equal to, less than, or greater than the specified value? See Workqueue-count Rule, page 9-40 .
Body Scanning	body-contains(<regular expression>)	Does the message contain text or an attachment that matches a specified pattern? Does the pattern occur the minimum number of times you specified for the threshold value? The engine scans delivery-status parts and associated attachments. See Body Scanning Rule, page 9-31 .
Body Scanning	only-body-contains(<regular expression>)	Does the message body contain text that matches a specified pattern? Does the pattern occur the minimum number of times you specified for the threshold value? Attachments are not scanned. See Body Scanning, page 9-31 .
Encryption Detection	encrypted	Is the message encrypted? See Encryption Detection Rule, page 9-32 .
Attachment Filename^a	attachment-filename	Does the message contain an attachment with a filename that matches a specific pattern? See Attachment Filename Rule, page 9-33 .
Attachment Type^a	attachment-type	Does the message contain an attachment of a particular MIME type? See Attachment Type Rule, page 9-33 .

Table 9-2 **Message Filter Rules**

Rule	Syntax	Description
Attachment File^a Type	attachment-filetype	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)? If the attachment is an Excel or Word document, you can also search for the following embedded file types: .exe , .dll, .bmp, .tiff, .pcx, .gif, .jpeg, png, and Photoshop images. You must enclose the file type in quotes to create a valid filter. You can use single or double quotes. For example, to search for .exe attachments, use the following syntax: <pre>if (attachment-filetype == "exe")</pre> For more information, see Examples of Attachment Scanning Message Filters, page 9-86.
Attachment MIME Type^a	attachment-mimetype	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.) See Examples of Attachment Scanning Message Filters, page 9-86.
Attachment Protected	attachment-protected	Does the message contain an attachment that is password protected? See Quarantining Protected Attachments, page 9-88.
Attachment Unprotected	attachment-unprotected	The attachment-unprotected filter condition returns true if the scanning engine detects an attachment that is unprotected. A file is considered unprotected if the scanning engine was able to read the attachment. A zip file is considered to be unprotected if any of its members is unprotected. Note — The attachment-unprotected filter condition is not mutually exclusive of the attachment-protected filter condition. It is possible for both filter conditions to return true when scanning the same attachment. This can occur, for example, if a zip file contains both protected and unprotected members. See Detecting Unprotected Attachments, page 9-89.

Table 9-2 Message Filter Rules

Rule	Syntax	Description
Attachment Scanning^a	<code>attachment-contains(<regular expression>)</code>	Does the message contain an attachment that contains text or another attachment that matches a specific pattern? Does the pattern occur the minimum number of times you specified for the threshold value? 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. See Examples of Attachment Scanning Message Filters, page 9-86.
Attachment Scanning	<code>attachment-binary-contains(<regular expression>)</code>	Does the message contain an attachment with binary data that matches a specific pattern? This rule is like the <code>attachment-contains()</code> rule, but it searches specifically for patterns in the binary data.
Attachment Scanning	<code>every-attachment-contains(<regular expression>)</code>	Do all of the attachments in this message contain text that matches a specific pattern? The text must exist in all attachments and the action performed is, in effect, a logical AND operation of an <code>'attachment-contains()'</code> for each attachment. The body is not scanned. Does the pattern occur the minimum number of times you specified for the threshold value? See Examples of Attachment Scanning Message Filters, page 9-86.
Attachment Size^a	<code>attachment-size</code>	Does the message contain an attachment whose size is within some range? This rule is similar to the <code>body-size</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. The size is evaluated prior to any decoding. See Examples of Attachment Scanning Message Filters, page 9-86.
Public Blacklists	<code>dnslist(<query server>)</code>	Does the sender’s IP address appear on a public blacklist server (RBL)? See DNS List Rule, page 9-34.
SenderBase Reputation	<code>reputation</code>	What is the sender’s SenderBase Reputation Score? See SenderBase Reputation Rule, page 9-35.
No SenderBase Reputation	<code>no-reputation</code>	Used to test if SenderBase Reputation Score is “None.” See SenderBase Reputation Rule, page 9-35.

Table 9-2 **Message Filter Rules**

Rule	Syntax	Description
Dictionary^b	<code>dictionary-match(<dictionary_name>)</code>	Does the message body contain any of the regular expressions or terms in the content dictionary named <i>dictionary_name</i> ? Does the pattern occur the minimum number of times you specified for the threshold value? See Dictionary Rules, page 9-36 .
Attachment Dictionary Match	<code>attachment-dictionary-match(<dictionary_name>)</code>	Does the attachment contain any of the regular expressions in the content dictionary named <i>dictionary_name</i> ? Does the pattern occur the minimum number of times you specified for the threshold value? See Dictionary Rules, page 9-36 .
Subject Dictionary Match	<code>subject-dictionary-match(<dictionary_name>)</code>	Does the Subject header contain any of the regular expressions or terms in the content dictionary named <i>dictionary name</i> ? See Dictionary Rules, page 9-36 .
Header Dictionary Match	<code>header-dictionary-match(<dictionary_name>, <header>)</code>	Does the specified header (case insensitive) contain any of the regular expressions or terms in the content dictionary named <i>dictionary name</i> ? See Dictionary Rules, page 9-36 .
Body Dictionary Match	<code>body-dictionary-match(<dictionary_name>)</code>	This filter condition returns true if the dictionary term matches content in the body of the message only. The filter searches for terms within the MIME parts not considered to be an attachment, and it returns true if the user-defined threshold is met (the default threshold value is one). See Dictionary Rules, page 9-36 .
Envelope Recipient Dictionary Match	<code>rcpt-to-dictionary-match(<dictionary_name>)</code>	Does the envelope recipient contain any of the regular expressions or terms in the content dictionary named <i>dictionary name</i> ? See Dictionary Rules, page 9-36 .
Envelope Sender Dictionary Match	<code>mail-from-dictionary-match(<dictionary_name>)</code>	Does the envelope sender contain any of the regular expressions or terms in the content dictionary named <i>dictionary name</i> ? See Dictionary Rules, page 9-36 .
SMTP Authenticated User Match	<code>smtp-auth-id-matches(<target> [, <sieve-char>])</code>	Does the address of the Envelope Sender and the address in message header match the authenticated SMTP user ID of the sender? See SMTP Authenticated User Match Rule, page 9-40 .
True	<code>true</code>	Matches all messages. See True Rule, page 9-23 .
Valid	<code>valid</code>	Returns false if the message contains unparsable/invalid MIME parts and true otherwise. See Valid Rule, page 9-23 .
Signed	<code>signed</code>	Is the message is signed? See Signed Rule, page 9-42 .

Table 9-2 **Message Filter Rules**

Rule	Syntax	Description
Signed Certificate	signed-certificate(<field> [<i><operator></i> <i><regular expression></i>])	Does the message signer or X.509 certificate issuer match a certain pattern? See Signed Certificate Rule , page 9-43.
Header Repeats	header-repeats (<target>, <threshold> [, <direction>])	Returns <i>true</i> if at a given point in time, a specified number of messages: • With same subject header are detected in last one hour. • From same envelope-sender are detected in last one hour. See Header Repeats Rule , page 9-45.
URL Reputation	url-reputation url-no-reputation	Is the reputation score of any URL in the message within the specified range? Is a reputation score for a URL unavailable? See URL Reputation Rules , page 9-47.
URL Category	url-category	Does the category of any URL in the message match the specified categories? See URL Category Rule , page 9-47.
Corrupt Attachment	attachment-corrupt	Does this message have an attachment that is corrupt? See Corrupt Attachment Rule , page 9-48.

a.Attachment filtering is discussed in detail in the section [Attachment Scanning](#), page 9-78.

b.Content Dictionaries are discussed in the detail in the “Text Resources” chapter.

Each message injected into the Cisco appliance is processed through all message filters in order, unless you specify a final action, which stops the message from being processed further. (See [Message Filter Actions](#), page 9-2.) Filters may also apply to all messages, and rules may also be combined using logical connectors (AND, OR, NOT).

Regular Expressions in Rules

Several of the atomic tests used to define rules use *regular expression matching*. Regular expressions can become complex. Use the following table as a guide for the applying of regular expressions within message filter rules:

Table 9-3 **Regular Expression in Rules**

Regular expression (abc)	Regular expressions in filter rules match a string if the sequence of directives in the regular expression match any part of the string. For example, the regular expression <code>Georg</code> matches the string <code>George Of The Jungle</code>, the string <code>Georgy Porgy</code>, the string <code>La Meson Georgette</code> as well as <code>Georg</code>.
Carat (^) Dollar sign (\$)	Rules containing the dollar sign character (\$) only match the end of the string, and rules containing the caret symbol (^) only match the beginning of the string. For example, the regular expression <code>^Georg\$</code> only matches the string <code>Georg</code>. Searching for an empty header would look like this: <code>"^\$"</code>
Letters, white space and the at sign (@) character	Rules containing characters, white space, and the at sign character (@) only match themselves explicitly. For example, the regular expression <code>^George@admin\$</code> only matches the string <code>George@admin</code>.
Period character (.)	Rules containing a period character (.) match any character (except a new line). For example, the regular expression <code>^...admin\$</code> matches the string <code>macadmin</code> as well as the string <code>sunadmin</code> but not <code>win32admin</code>.
Asterisk (*) directive	Rules containing an asterisk (*) match “zero or more matches of the previous directive.” In particular, the sequence of a period and an asterisk (.*) matches any sequence of characters (not containing a new line). For example, the regular expression <code>^P.*Piper\$</code> matches all of these strings: <code>PPiper</code>, <code>Peter Piper</code>, <code>P.Piper</code>, and <code>Penelope Penny Piper</code>.
Backslash special characters (\)	The backslash character <i>escapes</i> special characters. Thus the sequence <code>\.</code> only matches a literal period, the sequence <code>\\$</code> only matches a literal dollar sign, and the sequence <code>\^</code> only matches a literal caret symbol. For example, the regular expression <code>^ik\.ac\.uk\$</code> only matches the string <code>ik.ac.uk</code>. Important Note: The backslash is also a special escape character for the parser. As a result, if you want to include backslash in your regular expression, you must use <i>two</i> backslashes — so that after parsing, only one “real” backslash remains, which is then passed to the regular expression system. So, if you wanted to match the example domain above, you would enter <code>^ik\\.ac\\.uk\$</code>.

Table 9-3 Regular Expression in Rules

Case-insensitivity (<i>(?i)</i>)	The token <i>(?i)</i> that indicates the rest of the regular expression should be treated in case-insensitive mode. Placing this token at the beginning of a case-sensitive regular expression results in a completely insensitive match. For example, the regular expression “<i>(?i)viagra</i>” matches <i>Viagra</i>, <i>vIaGrA</i>, and <i>VIAGRA</i>.
Number of repetitions { <i>min,max</i> }	The regular expression notation that indicates the number of repetitions of the previous token is supported. For example, the expression “<i>fo{2,3}</i>” matches <i>foo</i> and <i>fooo</i> but not <i>fo</i> or <i>fofo</i>. This statement: <code>if(header('To') == "^.{500,}")</code> looks for a “To” header that has 500 or more characters in it.
Or (<i> </i>)	Alternation, or the “or” operator. If A and B are regular expressions, the expression “<i>A B</i>” will match any string that matches either “A” or “B.” For example, the expression “<i>foo bar</i>” will match either <i>foo</i> or <i>bar</i>, but not <i>foobar</i>.

Related Topics

- [Using Regular Expressions to Filter Messages, page 9-18](#)
- [Guidelines for Using Regular Expressions, page 9-19](#)
- [Regular Expression and Non-ASCII Character Sets, page 9-19](#)
- [n Tests, page 9-19](#)
- [Case-sensitivity, page 9-19](#)
- [Writing Efficient Filters, page 9-20](#)
- [PDFs and Regular Expressions, page 9-20](#)

Using Regular Expressions to Filter Messages

You can use filters to search for strings and patterns in non-ASCII encoded message content (both headers and bodies). Specifically, the system supports regular expression (regex) searching for non-ASCII character sets within:

- Message headers
- MIME attachment filename strings
- Message body:
 - Bodies without MIME headers (i.e. traditional email)
 - Bodies with MIME headers indicating encoding but no MIME parts
 - Multi-part MIME messages with encoding indicated
 - All of the above without the encoding specified in a MIME header

You can use regular expressions (regexes) to match on any part of the message or body, including matching attachments. The various attachment types include text, HTML, MS Word, Excel, and others. Examples of character sets of interest include gb2312, HZ, EUC, JIS, Shift-JIS, Big5, and Unicode.

Guidelines for Using Regular Expressions



<http://www.python.org/doc/howto/>

n Tests

```
rcpt-to != "^goober@dev\\.null\\.\\.\\.\\. $" (non-matching)
```

<https://t.me/learningnets>

Writing Efficient Filters

This example shows two filters that do the same thing, but the first one takes much more CPU. The second filter uses a regular expression that is more efficient.

```
attachment-filter: if ((recv-listener == "Inbound") AND
((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((
((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((
"\\.386$")) OR (attachment-filename == "\\..exe$")) OR (attachment-filename == "\\..ad$"))
OR (attachment-filename == "\\..ade$")) OR (attachment-filename == "\\..adp$")) OR
(attachment-filename == "\\..asp$")) OR (attachment-filename == "\\..bas$")) OR
(attachment-filename == "\\..bat$")) OR (attachment-filename == "\\..chm$")) OR
(attachment-filename == "\\..cmd$")) OR (attachment-filename == "\\..com$")) OR
(attachment-filename == "\\..cpl$")) OR (attachment-filename == "\\..crt$")) OR
(attachment-filename == "\\..exe$")) OR (attachment-filename == "\\..hlp$")) OR
(attachment-filename == "\\..hta$")) OR (attachment-filename == "\\..inf$")) OR
(attachment-filename == "\\..ins$")) OR (attachment-filename == "\\..isp$")) OR
(attachment-filename == "\\..js$")) OR (attachment-filename == "\\..jse$")) OR
(attachment-filename == "\\..lnk$")) OR (attachment-filename == "\\..mdb$")) OR
(attachment-filename == "\\..mde$")) OR (attachment-filename == "\\..msc$")) OR
(attachment-filename == "\\..msi$")) OR (attachment-filename == "\\..msp$")) OR
(attachment-filename == "\\..mst$")) OR (attachment-filename == "\\..pcd$")) OR
(attachment-filename == "\\..pif$")) OR (attachment-filename == "\\..reg$")) OR
(attachment-filename == "\\..scr$")) OR (attachment-filename == "\\..sct$")) OR
(attachment-filename == "\\..shb$")) OR (attachment-filename == "\\..shs$")) OR
(attachment-filename == "\\..url$")) OR (attachment-filename == "\\..vb$")) OR
(attachment-filename == "\\..vbe$")) OR (attachment-filename == "\\..vbs$")) OR
(attachment-filename == "\\..vss$")) OR (attachment-filename == "\\..vst$")) OR
(attachment-filename == "\\..vsw$")) OR (attachment-filename == "\\..ws$")) OR
(attachment-filename == "\\..wsc$")) OR (attachment-filename == "\\..wsf$")) OR
(attachment-filename == "\\..wsh$")) { bounce(); }
```

In this instance, AsyncOS will have to start the regular expression engine 30 times, once for each attachment type and the recv-listener.

Instead, write the filter to look like this:

```
attachment-filter: if (recv-listener == "Inbound") AND (attachment-filename ==
"\\. (386|exe|ad|ade|adp|asp|bas|bat|chm|cmd|com|cpl|crt|exe|hlp|hta|inf|ins|isp|js|jse|l
nk|mdb|mde|msc|msi|msp|mst|pcd|pif|reg|scr|sct|shb|shs|url|vb|vbe|vbs|vss|vst|vsw|ws|wsc
|wsf|wsh)$") {

    bounce();

}
```

The regular expression engine only has to start twice and the filter is arguably easier to maintain as you do not have to worry about adding “()”, spelling errors. In contrast to the above, this should show a decrease in CPU overhead.

PDFs and Regular Expressions

Depending on how a PDF is generated, it may contain no spaces or line breaks. When this occurs, the scanning engine attempts to insert logical spaces and line breaks based on the location of the words on the page. For example, when a word is constructed using multiple fonts or font sizes, the PDF code is

rendered in a way that makes it difficult for the scanning engine to determine word and line breaks. When you attempt to match a regular expression against a PDF file constructed in this way, the scanning engine may return unexpected results.

For example, you enter a word in a PowerPoint document that uses different fonts and different font sizes for each letter in the word. When a scanning engine reads a PDF generated from this application, it inserts logical spaces and line breaks. Because of the construction of the PDF, it may interpret the word “callout” as “call out” or “c a l l out.” Attempting to match either of these renderings against the regular expression, “callout,” would result in no matches.

Smart Identifiers

When you use message rules that scan message content, you can use smart identifiers to detect certain patterns in the data.

Smart identifiers can detect the following patterns in data:

- Credit card numbers
- U.S. Social Security numbers
- Committee on Uniform Security Identification Procedures (CUSIP) numbers
- American Banking Association (ABA) routing numbers

To use smart identifiers in a filter, enter the following keywords in a filter rule that scans body or attachment content:

Table 9-4 Smart Identifiers in Message Filters

Key Word	Smart Identifier	Description
*credit	Credit card number	Identifies 14-, 15-, and 16- digit credit card numbers. NOTE: The smart identifier does not identify enRoute or JCB cards.
*aba	ABA routing number	Identifies ABA routing numbers.
*ssn	Social security number	Identifies U.S. social security numbers. The *ssn smart identifier identifies social security numbers with dashes, periods and spaces.
*cusip	CUSIP number	Identifies CUSIP numbers.

Related Topics

- [Smart Identifier Syntax, page 9-21](#)

Smart Identifier Syntax

When you use a smart identifier in a filter rule, enter the smart-identifier keyword in quotes within a filter rule that scans the body or attachment file, as in the example below:

```
ID_Credit_Cards:
if (body-contains ("*credit")) {
```

```
notify("legaldept@example.com");  
  
}  
.
```

You can also use smart identifiers in content filters and as a part of content dictionaries.

**Note**

You cannot combine a smart identifier key word with a normal regular expression or another key word. For example the pattern `*credit|*ssn` would not be valid.

**Note**

To minimize on false positives using the `*SSN` smart identifier, it may be helpful to use the `*ssn` smart identifier along with other filter criteria. One example filter that can be used is the “only-body-contains” filter condition. This will only evaluate the expression to be true if the search string is present in all of the message body mime parts. For example, you could create the following filter:

```
SSN-nohtml: if only-body-contains("*ssn") { duplicate-quarantine("Policy"); }
```

Description and Examples of Message Filter Rules

The following section describes the various message filter rules in use and their examples.

Related Topics

- [True Rule, page 9-23](#)
- [Valid Rule, page 9-23](#)
- [Subject Rule, page 9-24](#)
- [Envelope Recipient Rule, page 9-24](#)
- [Envelope Recipient in Group Rule, page 9-25](#)
- [Envelope Sender Rule, page 9-25](#)
- [Envelope Sender in Group Rule, page 9-26](#)
- [Sender Group Rule, page 9-26](#)
- [Body Size Rule, page 9-27](#)
- [Remote IP Rule, page 9-27](#)
- [Receiving Listener Rule, page 9-28](#)
- [Receiving IP Interface Rule, page 9-28](#)
- [Date Rule, page 9-28](#)
- [Header Rule, page 9-29](#)
- [Random Rule, page 9-30](#)
- [Recipient Count Rule, page 9-30](#)
- [Address Count Rule, page 9-31](#)
- [Body Scanning Rule, page 9-31](#)

- [Body Scanning, page 9-31](#)
- [Encryption Detection Rule, page 9-32](#)
- [Attachment Type Rule, page 9-33](#)
- [Attachment Filename Rule, page 9-33](#)
- [DNS List Rule, page 9-34](#)
- [SenderBase Reputation Rule, page 9-35](#)
- [Dictionary Rules, page 9-36](#)
- [SPF-Status Rule, page 9-37](#)
- [SPF-Passed Rule, page 9-39](#)
- [S/MIME Gateway Message Rule, page 9-39](#)
- [S/MIME Gateway Verified Rule, page 9-40](#)
- [Workqueue-count Rule, page 9-40](#)
- [SMTP Authenticated User Match Rule, page 9-40](#)
- [Signed Rule, page 9-42](#)
- [Signed Certificate Rule, page 9-43](#)
- [Header Repeats Rule, page 9-45](#)
- [URL Reputation Rules, page 9-47](#)
- [URL Category Rule, page 9-47](#)
- [Corrupt Attachment Rule, page 9-48](#)

True Rule

The `true` rule matches all messages. For example, the following rule changes the IP interface to external for all messages it tests.

```
externalFilter:

    if (true)

    {

        alt-src-host('external');

    }
```

Valid Rule

The `valid` rule returns false if the message contains unparsable/invalid MIME parts and true otherwise. For example, the following rule drops all unparsable messages it tests.

```
not-valid-mime:

    if not valid

    {
```

```
        drop();
    }
}
```

Subject Rule

The `subject` rule selects those messages where the value of the subject header matches the given regular expression.

For example, the following filter discards all messages with subjects that start with the phrase `Make Money...`

```
scamFilter:

    if (subject == '^Make Money')

    {

        drop();

    }
```

You can specify non-ASCII characters to search for in the value of the header.

When working with headers, remember that the current value of the header includes changes made during processing (such as with filter actions that add, remove, or modify message headings). See [Message Header Rules and Evaluation, page 9-5](#) for more information.

The following filter returns true if the headers are empty or if the headers are missing from the message:

```
EmptySubject_To_filter:

if (header('Subject') != ".") OR

    (header('To') != ".") {

    drop();

}
```



Note

This filter returns true for empty Subject and To headers, but it also returns true for missing headers. If the message does not contain the specified headers, the filter still returns true.

Envelope Recipient Rule

The `rcpt-to` rule selects those messages where any Envelope Recipient matches the given regular expression. For example, the following filter drops all messages sent with an email address containing the string “scarface.”

**Note**

The regular expression for the `rcpt-to` rule is case *insensitive*.

```
scarfaceFilter:

    if (rcpt-to == 'scarface')

    {

        drop();

    }
```

**Note**

The `rcpt-to` rule is message-based. If a message has multiple recipients, only one recipient has to match the rule for the specified action to affect the message to all recipients.

Envelope Recipient in Group Rule

The `rcpt-to-group` rule selects those messages where any Envelope Recipient is found to be a member of the LDAP group given. For example, the following filter drops all messages sent with an email address within the LDAP group “ExpiredAccounts.”

```
expiredFilter:

    if (rcpt-to-group == 'ExpiredAccounts')

    {

        drop();

    }
```

**Note**

The `rcpt-to-group` rule is message-based. If a message has multiple recipients, only one recipient has to match the rule for the specified action to affect the message to all recipients.

Envelope Sender Rule

The `mail-from` rule selects those messages where the Envelope Sender matches the given regular expression. For example, the following filter immediately delivers any message sent by `admin@yourdomain.com`.

**Note**

The regular expression for the `mail-from` rule is case *insensitive*. Note that the period character is escaped in the following example.

```
kremFilter:

  if (mail-from == '^admin@yourdomain\\.com$')

  {

    skip-filters();

  }
```

Envelope Sender in Group Rule

The `mail-from-group` rule selects those messages where the Envelope Sender is found to be in an LDAP group given on the right side of the operator (or, in the case of inequality, where the sender's email address is *not* in the given LDAP group). For example, the following filter immediately delivers any message sent by someone whose email address is in the LDAP group "KnownSenders."

```
SenderLDAPGroupFilter:

  if (mail-from-group == 'KnownSenders')

  {

    skip-filters();

  }
```

Sender Group Rule

The `sendergroup` message filter selects a message based on which sender group was matched in a listener's Host Access Table (HAT). This rule uses '=' (for matching) or '!=' (for not matching) to test for matching a given regular expression (the right side of the expression). For example, the following message filter rule evaluates to `true` if the sender group of the message matches the regular expression `Internal`, and, if so, sends the message to an alternate mail host.

```
senderGroupFilter:

  if (sendergroup == "Internal")

  {

    alt-mailhost("[172.17.0.1]");

  }
```

Body Size Rule

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 given number. For example, the following filter bounces any message where the body size is greater than 5 megabytes.

```
BigFilter:

    if (body-size > 5M)

    {

        bounce();

    }
```

The `body-size` can be compared in the following ways:

Example	Comparison Type
<code>body-size < 10M</code>	Less than
<code>body-size <= 10M</code>	Less than or equal
<code>body-size > 10M</code>	Greater than
<code>body-size >= 10M</code>	Greater than or equal
<code>body-size == 10M</code>	Equal
<code>body-size != 10M</code>	Not equal

As a convenience, the size measurement may be specified with a suffix:

Quantity	Description
10b	ten bytes (same as 10)
13k	thirteen kilobytes
5M	five megabytes
40G	40 gigabytes (Note: The Cisco appliance cannot accept messages larger than 100 megabytes.)

Remote IP Rule

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 can be either Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6). The IP address pattern is specified using the **allowed hosts** notation described in “Sender Group Syntax”, except for the `SBO`, `SBRS`, `dnslist` notations and the special keyword `ALL`.

The allowed hosts notation can only identify sequences and numeric ranges of IP addresses (not hostnames). For example, the following filter bounces any message *not* injected from IP addresses of form `10.1.1.x` where `x` is 50, 51, 52, 53, 54, or 55.

```
notMineFilter:

    if (remote-ip != '10.1.1.50-55')
```

```

{
    bounce();
}

```

Receiving Listener Rule

The `recv-listener` rule selects those messages received on the named listener. The listener name must be the nickname of one of the listeners currently configured on the system. For example, the following filter immediately delivers any message arriving from the listener named `expedite`.

```

expediteFilter:

    if (recv-listener == 'expedite')

    {

        skip-filters();

    }

```

Receiving IP Interface Rule

The `recv-int` rule selects those messages received via the named interface. The interface name must be the nickname of one of the interfaces currently configured for the system. For example, the following filter bounces any message arriving from the interface named `outside`.

```

outsideFilter:

    if (recv-int == 'outside')

    {

        bounce();

    }

```

Date Rule

The `date` rule checks the current time and date against a time and date you specify. The date rule is compares against a string containing a timestamp of the format *MM/DD/YYYY hh:mm:ss*. This is useful to specify actions to be performed before or after certain times in US format. (Note that there may be an issue if you are searching messages with non-US date formats.) the following filter bounces all messages from `campaign1@yourdomain.com` that are injected after 1:00pm on July 28th, 2003:

```

TimeOutFilter:

    if ((date > '07/28/2003 13:00:00') and (mail-from ==

        'campaign1@yourdomain\\.com'))

```

```
{  
  
    bounce();  
  
}
```

**Note**

Do not confuse the `date` rule with the `$Date` message filter action variable.

Header Rule

The `header()` rule checks the message headers for a specific header, which must be specified quoted in parentheses (“*header name*”). This rule may be compared to a regular expression, much like the `subject` rule, or may be used without any comparison, in which case it will be “true” if the header is found in the message, and “false” if it is not found. For example, the following example checks to see if the header `X-Sample` is found, and if its value contains the string “`sample text`”. If a match is made, the message is bounced.

FooHeaderFilter:

```
if (header('X-Sample') == 'sample text')  
  
{  
  
    bounce();  
  
}
```

You can specify non-ASCII characters to search for in the value of the header.

The following example demonstrates the header rule without a comparison. In this case, if the header `X-DeleteMe` is found, it is removed from the message.

DeleteMeHeaderFilter:

```
if header('X-DeleteMe')  
  
{  
  
    strip-header('X-DeleteMe');  
  
}
```

When working with headers, remember that the current value of the header includes changes made during processing (such as with filter actions that add, remove, or modify message headings). See [Message Header Rules and Evaluation, page 9-5](#) for more information.

Random Rule

The `random` rule generates a random number from zero to N-1, where N is the integer value supplied in parenthesis after the rule. Like the `header()` rule, this rule may be used in a comparison, or may be used alone in a “unary” form. The rule evaluates to `true` in the unary form if the random number generated is non-zero. For example, both of the following filters are effectively equal, choosing Virtual Gateway address A half the time, and Virtual Gateway address B the other half of the time:

```
load_balance_a:

    if (random(10) < 5) {

        alt-src-host('interface_a');

    } else {

        alt-src-host('interface_b');

    }
```

```
load_balance_b:

    if (random(2)) {

        alt-src-host('interface_a');

    } else {

        alt-src-host('interface_b');

    }
```

Recipient Count Rule

The `rcpt-count` rule compares the number of recipients of a message against an integer value, in a similar way to the `body-size` rule. This can be useful for preventing users from sending email to large numbers of recipients at once, or for ensuring that such large mailing campaigns go out over a certain Virtual Gateway address. The following example sends any email with more than 100 recipients over a specific Virtual Gateway address:

```
large_list_filter:

    if (rcpt-count > 100) {

        alt-src-host('mass_mailing_interface');

    }
```

Address Count Rule

The `addr-count()` message filter rule takes one or more header strings, counts the number of recipients in each line and reports the cumulative number of recipients. This filter differs from the `rcpt-count` filter rule in that it operates on the message body headers instead of the envelope recipients. The following example shows the filter rule used to replace a long list of recipients with the alias, “undisclosed-recipients”:

```
count: if (addr-count("To", "Cc") > 30) {  
  
    strip-header("To");  
  
    strip-header("Cc");  
  
    insert-header("To", "undisclosed-recipients");  
  
}
```

Body Scanning Rule

The `body-contains()` rule scans the incoming email and all its attachments for a particular pattern defined by its parameter. This includes delivery-status parts and associated attachments. The `body-contains()` rule does not perform multi-line matching. The scanning logic can be modified on the Scan Behavior page or using the `scanconfig` command in the CLI to define which MIME types should or should not be scanned. You can also specify a minimum number of matches that the scanning engine must find in order for the scan to evaluate to true.

By default, the system scans all attachments *except* for those with a MIME type of `video/*`, `audio/*`, `image/*`. The system scans archive attachments — `.zip`, `.bzip`, `.compress`, `.tar`, or `.gzip` attachments containing multiple files. You can set the number of “nested” archived attachments to scan (for example, a `.zip` contained within a `.zip`).

For more information, see [Configuring Scan Behavior, page 9-115](#).

Body Scanning

When AsyncOS performs body scanning, it scans the body text and attachments for the regular expression. You can assign a minimum threshold value for the expression, and if the scanning engine encounters the regular expression the minimum number of times, the expression evaluates to true.

AsyncOS evaluates the different MIME parts of the message, and it scans any MIME part that is textual. AsyncOS identifies the text parts if the MIME type specifies text in the first part. AsyncOS determines the encoding based on the encoding specified in the message, and it converts the text to Unicode. It then searches for the regular expression in Unicode space. If no encoding is specified in the message, AsyncOS uses the encoding you specify on the Scan Behavior page or using the `scanconfig` command.

For more information about how AsyncOS evaluates MIME parts when scanning messages, see [Message Bodies vs. Message Attachments, page 9-5](#).

If the MIME part is not textual, AsyncOS extract files from a `.zip` or `.tar` archive or decompresses compressed files. After extracting the data, a scanning engine identifies the encoding for the file and returns the data from the file in Unicode. AsyncOS then searches for the regular expression in Unicode space.

The following example searches the body text and attachment for the phrase “Company Confidential.” The example specifies a minimum threshold of two instances, so if the scanning engine finds two or more instances of the phrase, it bounces any matching messages, and notifies the legal department of the attempt:

ConfidentialFilter:

```
if (body-contains('Company Confidential',2)) {

    notify ('legaldept@example.domain');

    bounce();

}
```

To scan only the body of the message, use `only-body-contains`:

disclaimer:

```
if (not only-body-contains('[dD]disclaimer',1) ) {

    notify('hresource@example.com');

}
```

Encryption Detection Rule

The `encrypted` rule examines the contents of a message for encrypted data. It does not attempt to decode the encrypted data, but merely examines the contents of the message for the existence of encrypted data. This can be useful for preventing users from sending encrypted email.



Note

The `encrypted` rule can only detect encrypted data in the content of messages. It does not detect encrypted attachments.

The `encrypted` rule is similar to the `true` rule in that it takes no parameters and cannot be compared. This rule returns `true` if encrypted data is found and `false` if no encrypted data is found. Because this function requires the message to be scanned, it uses the scanning settings you define on the Scan Behavior page or using the `scanconfig` command. For more information about configuring these options, see [Configuring Scan Behavior, page 9-115](#).

The following filter checks all email sent through the listener, and if a message contains encrypted data, the message is blind-carbon-copied to the legal department and then bounced:

prevent_encrypted_data:

```
if (encrypted) {

    bcc ('legaldept@example.domain');

    bounce();

}
```


Attachment Type Rule

The `attachment-type` rule checks the MIME types of each attachment in a message to see if it matches the given pattern. The pattern must be of the same form used in the Scan Behavior page or the `scanconfig` command, as described in [Configuring Scan Behavior, page 9-115](#), and so may have either side of the slash (/) replaced by an asterisk as a wildcard. If the message contains an attachment that matches this specified MIME type, this rule returns “true.”

Because this function requires the message to be scanned, it obeys all of the options described in [Configuring Scan Behavior, page 9-115](#).

See [Attachment Scanning, page 9-78](#) for more information on message filter rules you can use to manipulate attachments to messages.

The following filter checks all email sent through the listener, and if a message contains an attachment with a MIME type of `video/*`, the message is bounced:

```
bounce_video_clips:

    if (attachment-type == 'video/*') {

        bounce();

    }
```

Attachment Filename Rule

The `attachment-filename` rule checks the filenames of each attachment in a message to see if it matches the given regular expression. This comparison is case-sensitive. The comparison is, however sensitive to whitespace so if the filename has encoded whitespace at the end, the filter will skip the attachment. If one of the message’s attachments matches the filename, this rule returns “true.”

Please note the following points:

- Each attachment’s filename is captured from the MIME headers. The filename in the MIME header may contain trailing spaces.
- If an attachment is an archive, the Cisco appliance will harvest the filenames from inside the archive, and apply scan configuration rules (see [Configuring Scan Behavior, page 9-115](#)) accordingly.
 - If the attachment is a single compressed file (despite the file extension), it is not considered an archive and the filename of the compressed file is not harvested. This means that the file is not processed by the `attachment-filename` rule. An example of this type of file is an executable file (.exe) compressed with `gzip`.
 - For attachments consisting of a single compressed file, such as `foo.exe.gz`, use regular expression to search for specific file types within compressed files. See [Attachment Filenames and Single Compressed Files within Archive Files, page 9-34](#).

See [Attachment Scanning, page 9-78](#) for more information on message filter rules you can use to manipulate attachments to messages.

The following filter checks all email sent through the listener, and if a message contains an attachment with a filename `*.mp3`, the message is bounced:

```
block_mp3s:

    if (attachment-filename == '(?i)\\.mp3$') {
```

```

        bounce();
    }

```

Related Topics

- [Attachment Filenames and Single Compressed Files within Archive Files, page 9-34](#)

Attachment Filenames and Single Compressed Files within Archive Files

This example shows how to match single compressed files in archives such as those created by `gzip`:

```

quarantine_gzipped_exe_or_pif:

if (attachment-filename == '(?i)\\. (exe|pif) ($|.gz$)') {

    quarantine("Policy");

}

```

DNS List Rule

The `dnslist()` rule queries a public DNS List server that uses the DNSBL method (sometimes called “ip4r lookups”) of querying. The IP address of the incoming connection is reversed (so an IP of 1.2.3.4 becomes 4.3.2.1) and then added as a prefix to the server name in the parenthesis (a period to separate the two is added if the server name does not start with one). A DNS query is made, and the system is returned with either a DNS failure response (indicating the connection's IP address was not found in the server's list) or an IP address (indicating that the address was found). The IP address returned is *usually* of the form 127.0.0.x where x can be almost any number from 0 to 255 (IP address ranges are not allowed). Some servers actually return different numbers based on the reason for the listing, while others return the same result for all matches.

Like the `header()` rule, `dnslist()` can be used in either a unary or binary comparison. By itself, it simply evaluates to `true` if a response is received and `false` if no response is received (for example, if the DNS server is unreachable).

the following filter immediately delivers a message if the sender has been bonded with the Cisco Bonded Sender information services program:

```

whitelist_bondedsender:

    if (dnslist('query.bondedsender.org')) {

skip-filters();

    }

```

Optionally, you can compare the result to a string using the equality (`==`) or inequality (`!=`) expressions.

The following filter drops a message that results in a “127.0.0.2” response from the server. If the response is anything else, the rule returns “false” and the filter is ignored.

```
blacklist:

    if (dnslist('dnsbl.example.domain') == '127.0.0.2') {

        drop();

    }
```

SenderBase Reputation Rule

The `reputation` rule checks the SenderBase Reputation Score against another value. All the comparison operators are allowed, such as `>`, `==`, `<=`, and so forth. If the message does not have a SenderBase Reputation Score at all (because one was never checked for it, or because the system failed to get a response from the SenderBase Reputation Service query server), any comparison against a reputation fails (the number will not be greater than, less than, equal to, or not equal to any value). You can check for a SBRS score of “none” using the `no-reputation` rule described below. The following example adjusts the “Subject:” line of a message to be prefixed by “*** BadRep ***” if the reputation score returned from the SenderBase Reputation Service is below a threshold of -7.5..

```
note_bad_reps:

    if (reputation < -7.5) {

        strip-header ('Subject');

        insert-header ('Subject', '*** BadRep $Reputation *** $Subject');

    }
```

For more information, see the “Sender Reputation Filtering” chapter. See also [Bypass Anti-Spam System Action, page 9-72](#)

Values for the SenderBase Reputation rule are -10 through 10, but the value `NONE` may also be returned. To check specifically for the value `NONE`, use the `no-reputation` rule.

```
none_rep:

    if (no-reputation) {

        strip-header ('Subject');

        insert-header ('Subject', '*** Reputation = NONE *** $Subject');

    }
```

Dictionary Rules

The `dictionary-match(<dictionary_name>)` rule evaluates to `true` if the message body contains any of the regular expressions or terms in the content dictionary named “*dictionary_name*.” If the dictionary does not exist, the rule evaluates to `false`. For more information on defining dictionaries (including their case sensitivity and word boundary settings), see the “Text Resources” chapter.

The following filter blind carbon copies the administrator when the Cisco scans a message that contains any words within the dictionary named “secret_words.”

```
copy_codenames:

    if (dictionary-match ('secret_words')) {

        bcc('administrator@example.com');

    }
```

The following example sends the message to the Policy quarantine if the message body contains any words within the dictionary named “secret_words.” Unlike the `only-body-contains` condition, the `body-dictionary-match` condition does not require that all the content parts individually match the dictionary. The scores of each content part (taking into account multipart/alternative parts) are added together.

```
quarantine_data_loss_prevention:

    if (body-dictionary-match ('secret_words'))

    {

        quarantine('Policy');

    }
```

In the following filter, a subject that matches a term in the specified dictionary is quarantined:

```
quarantine_policy_subject:

    if (subject-dictionary-match ('gTest'))

    {

        quarantine('Policy');

    }
```

This example matches an email address in the “to” header and blind copies an administrator:

```
headerTest:

    if (header-dictionary-match ('competitorsList', 'to'))

    {
```

```
bcc('administrator@example.com');

}
```

The `attachment-dictionary-match(<dictionary_name>)` rule works like the `dictionary-match` rule above, except that it looks for matches in the attachment.

The following filter sends the message to the Policy quarantine if the message attachment contains any words found within the dictionary named “secret_words.”

```
quarantine_codenames_attachment:

if (attachment-dictionary-match ('secret_words'))

{

    quarantine('Policy');

}
```

The `header-dictionary-match(<dictionary_name>, <header>)` rule works like the `dictionary-match` rule above, except that it looks for matches in the header specified in `<header>`. The header name is case insensitive, so, for example, “subject” and “Subject” both work.

The following filter sends the message to the Policy quarantine if the message’s “cc” header contains any words found within the dictionary named “ex_employees.”

```
quarantine_codenames_attachment:

if (header-dictionary-match ('ex_employees', 'cc'))

{

    quarantine('Policy');

}
```

You can use wild cards within the dictionary terms. You do not have to escape the period in email addresses.

SPF-Status Rule

When you receive SPF/SIDF verified mail, you may want to take different actions depending on the results of the SPF/SIDF verification. The `spf-status` rule checks against different SPF verification results. For more information, see [Verification Results, page 20-31](#).

You can check against the SPF/SIDF verification results using the following syntax:

```
if (spf-status == "Pass")
```

If you want a single condition to check against multiple status verdicts, you can use the following syntax:

```
if (spf-status == "PermError, TempError")
```

You can also check the verification results against the HELO, MAIL FROM, and PRA identities using the following syntax:

```
if (spf-status("pra") == "Fail")
```

The following example shows the `spf-status` filter in use:

```
skip-spam-check-for-verified-senders:
```

```
    if (sendergroup == "TRUSTED" and spf-status == "Pass"){
        skip-spamcheck();
    }
```

```
quarantine-spf-failed-mail:
```

```
    if (spf-status("pra") == "Fail") {
        if (spf-status("mailfrom") == "Fail"){
            # completely malicious mail
            quarantine("Policy");
        } else {
            if(spf-status("mailfrom") == "SoftFail") {
                # malicious mail, but tempting
                quarantine("Policy");
            }
        }
    } else {
        if(spf-status("pra") == "SoftFail"){
            if (spf-status("mailfrom") == "Fail"
                or spf-status("mailfrom") == "SoftFail"){
                # malicious mail, but tempting
                quarantine("Policy");
            }
        }
    }
```

```

    }

    }

}

stamp-mail-with-spf-verification-error:

    if (spf-status("pra") == "PermError, TempError"

        or spf-status("mailfrom") == "PermError, TempError"

        or spf-status("helo") == "PermError, TempError"){

        # permanent error - stamp message subject

        strip-header("Subject");

        insert-header("Subject", "[POTENTIAL PHISHING] $Subject"); }

.

```

SPF-Passed Rule

The following example shows an `spf-passed` rule used to quarantine emails that are not marked as `spf-passed`:

```

quarantine-spf-unauthorized-mail:

    if (not spf-passed) {

        quarantine("Policy");

    }

```



Note

Unlike the `spf-status` rule, the `spf-passed` rule reduces the SPF/SIDF verification values to a simple Boolean. The following verification results are treated as not passed in the `spf-passed` rule: None, Neutral, Softfail, TempError, PermError, and Fail. To perform actions on messages based on more granular results, use the `spf-status` rule.

S/MIME Gateway Message Rule

The S/MIME Gateway Message rule checks if a message is S/MIME signed, encrypted, or signed and encrypted. The following message filter checks if the message is an S/MIME message and quarantines it if the verification or decryption using S/MIME fails.

```

quarantine_smime_messages:
if (smime-gateway-message and not smime-gateway-verified) {
quarantine("Policy");
}

```

For more information, see [Chapter 19, “S/MIME Security Services.”](#)

S/MIME Gateway Verified Rule

The S/MIME Gateway Message Verified rule checks if a message is successfully verified, decrypted, or decrypted and verified. The following message filter checks if the message is an S/MIME message and quarantines it if the verification or decryption using S/MIME fails.

```
quarantine_smime_messages:
if (smime-gateway-message and not smime-gateway-verified) {
quarantine("Policy");
}
```

For more information, see [Chapter 19, “S/MIME Security Services.”](#)

Workqueue-count Rule

The `workqueue-count` rule checks the `workqueue-count` against a specified value. All the comparison operators are allowed, such as `>`, `==`, `<=`, and so forth.

The following filter checks the workqueue count, and skips spam check if the queue is greater than the specified number.

```
wqfull:

if (workqueue-count > 1000) {

    skip-spamcheck();

}
```

For more information on SPF/SIDF, see [Overview of SPF and SIDF Verification, page 20-22.](#)

SMTP Authenticated User Match Rule

If your Cisco appliance uses SMTP authentication to send messages, the `smtp-auth-id-matches (<target> [, <sieve-char>])` rule can check a message’s headers and Envelope Sender against the sender’s SMTP authenticated user ID to identify outgoing messages with spoofed headers. This filter allows the system to quarantine or block potentially spoofed messages.

The `smtp-auth-id-matches` rule compares the SMTP authenticated ID against the following targets:

Target	Description
*EnvelopeFrom	Compares the address of the Envelope Sender (also known as MAIL FROM) in the SMTP conversation
*FromAddress	Compares the addresses parsed out of the From header. Since multiple addresses are permitted in the From: header, only one has to match.
*Sender	Compares the address specified in the Sender header.

Target	Description
*Any	Matches messages that were created during an authenticated SMTP session regardless of identity.
*None	Matches messages that were not created during an authenticated SMTP session. This is useful when authentication is optional (preferred).

The filter performs matches loosely. It is not case-sensitive. If the optional *sieve-char* parameter is supplied, the last portion of an address that follows the specified character will be ignored for the purposes of comparison. For example, if the + character is included as a parameter, the filter ignores the portion of the address `joe+folder@example.com` that follows the + character. If the address was `joe+smith+folder@example.com`, only the `+folder` portion is ignored. If the SMTP authenticated user ID string is a simple username and not a fully-qualified e-mail address, only the username portion of the target will be examined to determine a match. The domain must be verified in a separate rule.

Also, you can use the `$SMTPAuthID` variable to insert the SMTP authenticated user ID into headers.

The following table shows examples of comparisons between the SMTP authenticated ID and email addresses and whether they would match using the `smtp-auth-id-matches` filter rule:

SMTP Auth ID	Sieve Char	Comparison Address	Matches?
someuser		otheruser@example.com	No
someuser		someuser@example.com	Yes
someuser		someuser@another.com	Yes
SomeUser		someuser@example.com	Yes
someuser		someuser+folder@example.com	No
someuser	+	someuser+folder@example.com	Yes
someuser@example.com		someuser@forged.com	No
someuser@example.com		someuser@example.com	Yes
SomeUser@example.com		someuser@example.com	Yes

The following filter checks all messages created during an authenticated SMTP session to verify that the addresses in the From header and the Envelope Sender match the SMTP authenticated user ID. If the addresses and the ID match, the filter verifies the domain. If they do not match, the appliance quarantines the message.

```
Msg_Authentication:
```

```
if (smtp-auth-id-matches ("*Any"))
```

```
{
```

```
    # Always include the original authentication credentials in a
```

```
    # special header.
```

```
    insert-header ("X-Auth-ID", "$SMTPAuthID");
```

```

if (smtp-auth-id-matches("*FromAddress", "+") and
    smtp-auth-id-matches("*EnvelopeFrom", "+"))
{
    # Username matches. Verify the domain
    if header('from') != "(?i)@(:example\\.com|alternate\\.com)" or
        mail-from != "(?i)@(:example\\.com|alternate\\.com)"
    {
        # User has specified a domain which cannot be authenticated
        quarantine("forged");
    }
} else {
    # User claims to be an completely different user
    quarantine("forged");
}
}

```

Signed Rule

The `signed` rule checks messages for a signature. The rule returns a boolean value to indicate if the message is signed or not. This rule evaluates whether the signature is encoded according to ASN.1 DER encoding rules and that it conforms to the CMS SignedData Type structure (RFC 3852, Section 5.1.). It does not aim to validate whether the signature matches the content, nor does it check the validity of the certificate.

The following example shows a `signed` rule used to insert headers into a signed message:

```
signedcheck: if signed { insert-header("X-Signed", "True"); }
```

The following example shows a `signed` rule used to drop attachments from unsigned messages from a certain sender group:

```

Signed: if ((sendergroup == "NOTTRUSTED") AND NOT signed) {
    html-convert();

    if (attachment_size > 0)
    {
        drop_attachments("");
    }
}

```

```

    }
}

```

Signed Certificate Rule

The `signed-certificate` rule selects those S/MIME messages where the X.509 certificate issuer or message signer matches the given regular expression. This rule only supports X.509 certificates.

The rule's syntax is `signed-certificate (<field> [<operator> <regular expression>])`, where:

- `<field>` is either the quoted string "issuer" or "signer",
- `<operator>` is either `==` or `!=`,
- and `<regular expression>` is the value for matching the "issuer" or "signer."

If the message is signed using multiple signatures, the rule returns true if any of the issuers or signers match the regular expression. The short form of this rule, `signed-certificate("issuer")` and `signed-certificate("signer")`, returns true if the S/MIME message contains an issuer or signer.

Related Topics

- [Signer, page 9-43](#)
- [Issuer, page 9-43](#)
- [Escaping in Regular Expressions, page 9-43](#)
- [\\$CertificateSigners Action Variable, page 9-44](#)
- [Examples, page 9-45](#)

Signer

For message signers, the rule extracts the sequence of `rfc822Name` names from the X.509 certificate's `subjectAltName` extension. If there is no `subjectAltName` field in the signing certificate, or this field does not have any `rfc822Name` names, the `signed-certificate("signer")` rule evaluates to false. In the rare cases of multiple `rfc822Name` names, the rule tries to match all of the names to the regular expression and evaluates as true on the first match.

Issuer

The issuer is a non-empty distinguished name in the X.509 certificate. AsyncOS extracts the issuer from the certificate and converts it to an LDAP-UTF8 Unicode string. For example:

- `C=US,S=CA,O=IronPort`
- `C=US,CN=Bob Smith`

Since X.509 certificates require the issuer field, `signed-certificate("issuer")` evaluates whether the S/MIME message contains an X.509 certificate.

Escaping in Regular Expressions

LDAP-UTF8 defines a mechanism for escaping that you can use in your regular expressions. For a detailed discussion on escaping characters in LDAP-UTF8, consult Lightweight Directory Access Protocol (LDAP): String Representation of Distinguished Names, accessible from <http://www.ietf.org/rfc/rfc4514.txt>.

The escaping rules for the `signed-certificate` rule's regular expressions differ from the escaping rules defined in LDAP-UTF8 by limiting escaping to only the characters that require escaping. LDAP-UTF8 allows optional escaping for characters that can be represented without escaping. For example, the following two strings are considered correct for "Example, Inc." using the LDAP-UTF8 escaping rules:

- `Example\, Inc.`
- `Example\\,\\ Inc\\.`

However, the `signed-certificate` rule only matches `Example\, Inc.` The regular expression does not allow escaping the space and period for matching because these characters do not require escaping, even though it is permitted in LDAP-UTF8. When creating a regular expression for the `signed-certificate` rule, do not escape a character if it can be represented without escaping.

\$CertificateSigners Action Variable

The action variable `$CertificateSigners` is a comma separated list of signers obtained from the `subjectAltName` element of the signing certificate. Multiple email addresses of a single signer will be included in the list with duplicates removed.

For example, Alice signs a message with her two certificates. Bob signs the message with his single certificate. All certificates are issued by a single corporate authority. After the message passes the S/MIME scan, the extracted data contain three items:

```
[
  {
    'issuer': 'CN=Auth,O=Example\, Inc.',
    'signer': ['alice@example.com', 'al@private.example.com']
  },
  {
    'issuer': 'CN=Auth,O=Example\, Inc.',
    'signer': ['alice@example.com', 'al@private.example.com']
  },
  {
    'issuer': 'CN=Auth,O=Example\, Inc.',
    'signer': ['bob@example.com', 'bob@private.example.com']
  }
]
```

The `$CertificateSigners` variable expands to:

```
"alice@example.com, al@private.example.com, bob@example.com, bob@private.example.com"
```

Examples

The following example inserts a new header if the certificate issuer is from the US:

```
Issuer: if signed-certificate("issuer") == "(?i)C=US" {
    insert-header("X-Test", "US issuer");
}
```

The following example notifies an administrator if the signer is not from example.com:

```
NotOurSigners: if signed-certificate("signer") AND
    signed-certificate("signer") != "example\\.com$" {
    notify("admin@example.com");
}
```

The following example adds a header if the message has an X.509 certificate:

```
AnyX509: if signed-certificate ("issuer") {
    insert-header("X-Test", "X.509 present");
}
```

The following example adds a header if the message's certificate does not have a signer:

```
NoSigner: if not signed-certificate ("signer") {
    insert-header("X-Test", "Old X.509?");
}
```

Header Repeats Rule

The Header Repeats rule evaluates to `true` if at a given point in time, a specified number of messages:

- With same subject are detected in the last one hour.
- From same envelope sender are detected in the last one hour.

You can use this rule to detect high volume emails. For example, political campaigns through certain websites may send out emails to organizations in high volumes. Anti-spam engines treat such emails as clean, and do not stop the delivery of these emails.

The syntax of this rule is `header-repeats (<target>, <threshold> [, <direction>])`, where:

- `<target>` is `subject` or `mail-from`. AsyncOS counts the repetition of values of the target.
- `<threshold>` is the number of messages with identical values for a given target, received in the last one hour, beyond which the rule evaluates to `true`.

- `<direction>` is `incoming`, `outgoing`, or `both`. If direction is not specified in this rule, incoming or outgoing messages are counted for rule evaluation.

Every time when a Header Repeats rule evaluates to `true`, a System Alert is sent. See [System Alerts, page 33-42](#).

**Note**

If the header field includes comma or semi-colon separated values, the rule considers the complete string for tracking. This rule ignores messages with empty subject header.

The Header Repeats rule maintains a moving sum of messages with up to one minute's precision. As a result, after the set threshold has reached, there can be a delay of one minute before this rule is triggered.

Related Topics

- [Using Header Repeats Rule with Other Rules, page 9-46](#)
- [Examples, page 9-46](#)

Using Header Repeats Rule with Other Rules

You can use the Header Repeats rule with other rules using AND or OR operators. For example, you can whitelist a subset of messages using the following filter:

```
F1: if (recv_listener == 'Gray') AND (header-repeats('subject', X, 'incoming')) {
drop();}
```

When you use a Header Repeats rule with another rule using AND or OR operators, the Header Repeats rule is evaluated last, and only if needed. If a Header Repeats rule is not evaluated for a given message, `subject` or `mail-from` is not counted to compare with the supplied threshold.

As Header Repeats rule is evaluated last and only if needed, the behavior of this rule may vary when used with other rules using an OR operator. The following sample filter uses an OR condition of Signed and Header Repeats rule.

```
f1: if signed OR (header-repeats('subject', 10)) { drop();}
```

In this example, if the first nine messages processed by this filter are signed messages with identical subject, the Header Repeats rule will not process these messages. If the tenth message is an unsigned message with identical subject header as the previous nine messages, the filter will not perform the configured action, even though the threshold has reached.

Examples

In the following example, at any given point in time, if the filter detects `x` or more incoming messages with identical subject in the last one hour, the subsequent messages with identical subject are sent to Policy quarantine.

```
f1 : if header-repeats('subject', X, 'incoming') { quarantine('Policy');}
```

In the following example, at any given point in time, if the filter detects *x* or more outgoing messages from same envelope sender in the last one hour, the subsequent messages from the same envelope sender are dropped and discarded.

```
f2 : if header-repeats('mail-from', X, 'outgoing') {drop();}
```

In the following example, at any given point in time, if the filter detects *x* or more incoming or outgoing messages with identical subject in the last one hour, the administrator is notified for every subsequent message with identical subject.

```
f3: if header-repeats('subject', X) {notify('admin@xyz.com');}
```

URL Reputation Rules

Use a URL reputation rule to define message actions based on the reputation score of any URL in the message. For important details, see [Filtering by URL Reputation or URL Category: Conditions and Rules, page 15-8 in Chapter 15, “URL Filtering.”](#)

For these rules:

- `msg_filter_name`: is the name of this message filter.
- `whitelist` is the name of a defined URL list (via the `urllistconfig` command.) Specifying a `whitelist` is optional.

To take action when the reputation service provides a score:

Use the `url-reputation` rule.

Filter syntax when using a `url-reputation` rule is:

```
<msg_filter_name>:
if url-reputation(<min_score>, <max_score>, '<whitelist>')
{<action>}
```

Where:

- `min_score` and `max_score` are the minimum and maximum scores in the range for which the action should apply. The values that you specify are included in the range.

Minimum and maximum scores must be between -10.0 and 10.0.

To take action when the reputation service does not provide a score:

Use the `url-no-reputation` rule.

Filter syntax when using a `url-no-reputation` rule is:

```
<msg_filter_name>:
if url_no_reputation('<whitelist>')
{<action>}
```

URL Category Rule

Use URL categories to define message actions based on the category of URLs in the message. For important details, see [Filtering by URL Reputation or URL Category: Conditions and Rules, page 15-8 in Chapter 15, “URL Filtering.”](#)

Filter syntax when using a url-category rule is:

```
<msg_filter_name>: if url-category ([ '<category-name1>', '<category-name2>', ...,
'<category-name3>' ], '<url_white_list>')
{
<action>
}
```

Where:

- `msg_filter_name` is the name of this message filter.
- `action` is any Message Filter action.
- `category-name` is the URL category. Separate multiple categories with commas. To obtain correct category names, look at a URL Category condition or action in a Content Filter. For descriptions and examples of the categories, see [About URL Categories, page 15-13](#).
- `url_white_list` is the name of a defined URL list (via the `urllistconfig` command.)

Corrupt Attachment Rule

The Corrupt Attachment rule evaluates to `true` if a message contains corrupt attachment. A corrupt attachment is an attachment that the scanning engine cannot scan and identified as corrupt.

Related Topics

- [Example, page 9-48](#)

Example

In the following example, if the filter detects a corrupt attachment in a message, the message is quarantined to Policy Quarantine.

```
quar_corrupt_attach: if (attachment-corrupt) { quarantine("Policy"); }
```


Message Filter Actions

The purpose of message filters is to perform actions on selected messages.

The two types of actions are:

- *Final* actions — such as `deliver`, `drop`, and `bounce` — end the processing of a message, and permit no further processing through subsequent filters.
- *Non-final* actions perform an action which permits the message to be processed further.

Non-final message filter actions are cumulative. If a message matches multiple filters where each filter specifies a different action, then all actions are accumulated and enforced. However, if a message matches multiple filters specifying the same action, the prior actions are overridden and the final filter action is enforced.

Related Topics

- [Filter Actions Summary Table, page 9-49](#)
- [Action Variables, page 9-56](#)
- [Matched Content Visibility, page 9-58](#)
- [Description and Examples of Message Filter Actions, page 9-58](#)

Filter Actions Summary Table

Message filters can apply the following actions shown in [Table 9-5](#) to an email message.

Table 9-5 *Message Filter Actions*

Action	Syntax	Description
Alter source host	<code>alt-src-host</code>	Change the source hostname and IP interface (Virtual Gateway address) to send the message. See Alter Source Host (Virtual Gateway address) Action, page 9-68 .
Alter recipient	<code>alt-rcpt-to</code>	Change a recipient of the message. See Alter Recipient Action, page 9-66 .
Alter mailhost	<code>alt-mailhost</code>	Change the destination mail host for the message. See Alter Delivery Host Action, page 9-67 .
Notify	<code>notify</code>	Report this message to another recipient. See Notify and Notify-Copy Actions, page 9-61 .
Notify Copy	<code>notify-copy</code>	Perform just like the <code>notify</code> action, but also sends a copy as with the <code>bcc-scan</code> action. See Notify and Notify-Copy Actions, page 9-61 .
Blind carbon copy	<code>bcc</code>	Copy this message (message replication) anonymously to another recipient. See Blind Carbon Copy Actions, page 9-63 .
Blind carbon copy with scan	<code>bcc-scan</code>	Copy this message anonymously to another recipient, and process that message through the work queue as if it were a new message. See Blind Carbon Copy Actions, page 9-63 .

Table 9-5 **Message Filter Actions**

Action	Syntax	Description
Archive	archive	Archive this message into an mbox-format file. See Archive Action, page 9-68 .
Quarantine	quarantine (<i>quarantine_name</i>)	Flag this message to be sent to the quarantine named <i>quarantine_name</i> . See Quarantine and Duplicate Actions, page 9-65 .
Duplicate (Quarantine)	duplicate-quarantine(<i>quarantine_name</i>)	Send a copy of the message to the specified quarantine. See Quarantine and Duplicate Actions, page 9-65 .
Remove headers	strip-header	Remove specified headers from the message before delivering. See Strip Header Action, page 9-69 .
Insert headers	insert-header	Insert a header and value pair into the message before delivering. See Insert Header Action, page 9-69 .
Edit header text	edit-header-text	Replace specified header text with a text string you specify in the filter condition. See Edit Header Text Action, page 9-70 .
Edit body text	edit-body-text()	Strip a regular expression from a message body and replaces it with text that you specify. You might want to use this filter if you want to remove and replace specific content, such as a URL within a message body. See Edit Body Text Action, page 9-70 .
Convert HTML	html-convert()	Strip HTML tags from message bodies and leaves the plain text content of the message. You might want to use this filter if you want to convert all HTML text in a message to plain text. HTML Convert Action, page 9-71 .
Assign bounce profile	bounce-profile	Assign a specific bounce profile to the message. See Bounce Profile Action, page 9-72 .
Bypass Anti-Spam System	skip-spamcheck	Ensure that the anti-spam systems in the Cisco system are <i>not</i> applied to this message. See Bypass Anti-Spam System Action, page 9-72 .
Bypass Anti-Virus System	skip-viruscheck	Ensure that the anti-virus systems in the Cisco system are <i>not</i> applied to this message. See Bypass Anti-Virus System Action, page 9-73 .
Bypass File Reputation Filtering and File Analysis	skip-ampcheck	Ensure that File Reputation Filtering and File Analysis are <i>not</i> applied to this message. See Bypass File Reputation Filtering and File Analysis System Actions, page 9-73 .
Skip Outbreak Filter Scanning	skip-vofcheck	Ensure that this message is not processed by the Outbreak Filters scanning. See Bypass Anti-Virus System Action, page 9-73 .
Drop Attachments by Name	drop-attachments-by-name	Drop 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. See Examples of Attachment Scanning Message Filters, page 9-86 .

Table 9-5 **Message Filter Actions**

Action	Syntax	Description
Drop Attachments by Type	<code>drop-attachments-by-type</code>	Drop all attachments on messages that have a MIME type, determined by either the given MIME type or the file extension. Archive file attachments (zip, tar) will be dropped if they contain a file that matches. See Examples of Attachment Scanning Message Filters , page 9-86.
Drop Attachments by File Type	<code>drop-attachments-by-filetype</code>	Drop 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 Scanning , page 9-78.
Drop Attachments by MIME Type	<code>drop-attachments-by-mimetype</code>	Drop all attachments on messages that have a given MIME type. This action does not attempt to ascertain the MIME type by file extension and so it also does not examine the contents of archives. See Examples of Attachment Scanning Message Filters , page 9-86.
Drop Attachments by Size	<code>drop-attachments-by-size</code>	Drop 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 prior to any decoding. See Examples of Attachment Scanning Message Filters , page 9-86.
Drop Attachments by Content	<code>drop-attachments-where-contains</code>	Drop all attachments on message that contain the regular expression. Does the pattern occur the minimum number of times you specified for the threshold value? Archive files (zip, tar) will be dropped if any of the files they contain match the regular expression pattern. See Examples of Attachment Scanning Message Filters , page 9-86. 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.
Drop Attachments by Dictionary Matches	<code>drop-attachments-where-dictionary-match</code>	Strip attachments based on matches to dictionary terms. If the terms in the MIME parts considered to be an attachment match a dictionary term (and the user-defined threshold is met), the attachment is stripped from the email. See Examples of Attachment Scanning Message Filters , page 9-86.
Add Footer	<code>add-footer(<i>footer-name</i>)</code>	Add disclaimer text as a footer to the message. See “Message Disclaimer Stamping” in the “Text Resources” chapter for more information.
Add Heading	<code>add-heading(<i>heading-name</i>)</code>	Add disclaimer text as a heading to the message. See “Message Disclaimer Stamping” in the “Text Resources” chapter for more information.

Table 9-5 **Message Filter Actions**

Action	Syntax	Description
Encrypt on Delivery	encrypt-deferred	Encrypt message on delivery, which means that the message continues to the next stage of processing, and when all processing is complete, the message is encrypted and delivered.
S/MIME Sign/Encrypt on Delivery	smime-gateway-deferred ("sending_profile")	Performs an S/MIME signing or encryption of the message using the specified sending profile during the delivery. See S/MIME Sign or Encrypt on Delivery Action , page 9-61.
S/MIME Sign/Encrypt	smime-gateway("sending_profile")	Performs an S/MIME signing or encryption using the specified sending profile and delivers the message, skipping any further processing. See S/MIME Sign or Encrypt Action , page 9-61.
Add Message Tag	tag-message(tag-name)	Add a custom term into the message to use with RSA Email DLP policy filtering. You can configure a RSA Email DLP policy to limit scanning to messages with the message tag. The message tag is not visible to recipients. See Add Message Tag Action , page 9-74 and the “Data Loss Prevention” chapter.
Add Log Entry	log-entry	Adds customized text into the Text Mail logs at the INFO level. The text can include action variables. The log entry appears in message tracking. See Add Log Entry Action , page 9-74.
Replace URL with text, based on URL reputation	- url-reputation-replace - url-no-reputation-replace	Modify URLs or their behavior based on the reputation of the URL. Use a separate action to handle the case in which the reputation service does not provide a score for a URL. See URL Reputation Actions , page 9-75.
Defang URL based on URL reputation	- url-reputation-defang - url-no-reputation-defang	
Redirect URL to a Cisco security proxy, based on URL reputation	- url-reputation-proxy-redirect - url-no-reputation-proxy-redirect	
Replace URL with text, based on URL Category	url-category-replace	Modify URLs or their behavior based on the category of the URL. See URL Category Actions , page 9-76.
Defang URL based on URL category	url-category-defang	
Redirect URL to Cisco security proxy, based on URL category	url-category-proxy-redirect	
No Operation	no-op	No action is performed. See No Operation , page 9-78.

Table 9-5 *Message Filter Actions*

Action	Syntax	Description
*Skip Remaining Message Filters	skip-filters	Ensure that this message is not processed by any other message filters and continues through the email pipeline. See Skip Remaining Message Filters Action, page 9-59 .
*Drop message	drop	Drop and discard the message. See Drop Action, page 9-59 .
*Bounce message	bounce	Send the message back to the sender. See Bounce Action, page 9-60 .
*Encrypt and Deliver Now	encrypt	Use Cisco Email Encryption to encrypt outgoing messages. See Encrypt Action, page 9-60 .
* Final Actions		

Related Topics

- [Attachment Groups, page 9-53](#)

Attachment Groups

You can specify a particular file type (“exe” files for example) or common groups of attachments in the `attachment-filetype` and `drop-attachments-by-filetype` rules. AsyncOS divides the attachments into the groups listed in [Table 9-6](#).

If you create a message filter that uses the `!=` operator to match a message that does not contain an attachment with a specific file type, the filter will not perform any action on the message if there is at least one attachment with the file type you want to filter out. For example, the following filter drops any message with an attachment that is not an `.exe` file type:

```
exe_check: if (attachment-filetype != "exe") {
    drop();
}
```

If a message has multiple attachments, the Email Security appliance does not drop the message if at least one of the attachments is an `.exe` file, even if the other attachments not `.exe` files.

Table 9-6 Attachment Groups

Attachment Group Name	Scanned File Types
Document	• doc • docx • mdb • mpp • ole • pdf • ppt • pptx • rtf • wps • x-wmf • xls • xlsx
Executable	• exe • java • msi • pif Note Filtering the Executable group will also scan .dll and .scr files, but you cannot filter these file types individually.
Compressed	• ace (ACE Archiver compressed file) • arc (SQUASH Compressed archive) • arj (Robert Jung ARJ compressed archive) • binhex • bz (Bzip compressed file) • bz2 (Bzip compressed file) • cab (Microsoft cabinet file) • gzip* (Compressed file - UNIX gzip) • lha (Compressed Archive [LHA/LHARC/LHZ]) • rar (Compressed archive) • sit (Compressed archive - Macintosh file [Stuffit]) • tar* (Compressed archive) • unix (UNIX compress file) • zip* (Compressed archive - Windows) • zoo (ZOO Compressed Archive File) * These file types can be “body-scanned”

Table 9-6 Attachment Groups (continued)

Attachment Group Name	Scanned File Types
Text	• txt • html • xml
Image	• bmp • cur • gif • ico • jpeg • pcx • png • psd • psp • tga • tiff
Media	• aac • aiff • asf • avi • flash • midi • mov • mp3 • mpeg • ogg • ram • snd • wav • wma • wmv

Action Variables

The `bcc()`, `bcc-scan()`, `notify()`, `notify-copy()`, `add-footer()`, `add-heading()`, and `insert-headers()` actions have parameters that may use certain variables that will be automatically replaced with information from the original message when the action is executed. These special variables are called *action variables*. Your Cisco appliance supports the following set of action variables:

Table 9-7 Message Filter Action Variables

Variable	Syntax	Description
All Headers	<code>\$AllHeaders</code>	Returns the message headers.
Body Size	<code>\$BodySize</code>	Returns the size, in bytes, of the message.
Certificate Signers	<code>\$CertificateSigners</code>	Returns the signers from the <code>subjectAltName</code> element of a signing certificate. See \$CertificateSigners Action Variable, page 9-44 for more information.
Date	<code>\$Date</code>	Returns the current date, using the format MM/DD/YYYY.
Dropped File Name	<code>\$dropped_filename</code>	Returns only the most recently dropped filename.
Dropped File Names	<code>\$dropped_filenames</code>	Displays list of dropped files (similar to <code>\$filenames</code>).
Dropped File Types	<code>\$dropped_filetypes</code>	Displays list of dropped file types (similar to <code>\$filetypes</code>).
Envelope Sender	<code>\$EnvelopeFrom</code>	Returns the Envelope Sender (Envelope From, <MAIL FROM>) of the message.
Envelope Recipients	<code>\$EnvelopeRecipients</code>	Returns all Envelope Recipients (Envelope To, <RCPT TO>) of the message.
File Names	<code>\$filenames</code>	Returns a comma-separated list of the message's attachments' filenames.
File Sizes	<code>\$filesizes</code>	Returns a comma-separated list of the message's attachments file sizes.
File Types	<code>\$filetypes</code>	Returns a comma-separated list of the message's attachments' file types.
Filter Name	<code>\$FilterName</code>	Returns the name of the filter being processed.
GMTTimeStamp	<code>\$GMTTimeStamp</code>	Returns the current time and date, as would be found in the Received: line of an email message, using GMT.
HAT Group Name	<code>\$Group</code>	Returns 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.
Matched Content	<code>\$MatchedContent</code>	Returns the content that triggered a scanning filter rule (including filter rules such as <code>body-contains</code> and content dictionaries).

Table 9-7 *Message Filter Action Variables (continued)*

Variable	Syntax	Description
Mail Flow Policy	<code>\$Policy</code>	Returns 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.
Header	<code>\$Header['string']</code>	Returns the value of the quoted header, if the original message contains a matching header. Note that double quotes may also be used.
Hostname	<code>\$Hostname</code>	Returns the hostname of the Cisco appliance.
Internal Message ID	<code>\$MID</code>	Returns the Message ID, or "MID" used internally to identify the message. Not to be confused with the RFC822 "Message-Id" value (use <code>\$Header</code> to retrieve that).
Receiving Listener	<code>\$RecvListener</code>	Replaced by the nickname of the listener that received the message.
Receiving Interface	<code>\$RecvInt</code>	Returns the nickname of the interface that received the message.
Remote IP Address	<code>\$RemoteIP</code>	Returns the IP address of the system that sent the message to the Cisco appliance.
Remote Host Address	<code>\$remotehost</code>	Returns the hostname of the system that sent the message to the Cisco appliance.
SenderBase Reputation Score	<code>\$Reputation</code>	Returns the SenderBase Reputation score of the sender. If there is no reputation score, it is replaced with "None".
Subject	<code>\$Subject</code>	Returns the subject of the message.
Time	<code>\$Time</code>	Returns the current time, in the local time zone.
Timestamp	<code>\$Timestamp</code>	Returns the current time and date, as would be found in the Received: line of an email message, in the local time zone.

Related Topics

- [Non-ASCII Character Sets and Message Filter Action Variables, page 9-57](#)

Non-ASCII Character Sets and Message Filter Action Variables


The system supports the expansion of action variables that contain ISO-2022 style character codings (the style of encoding used in header values) and also supports international text in the notification. These will be merged together to generate a notification that will then be sent as a UTF-8, quoted printable message.

Matched Content Visibility

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.

When you view messages in the local quarantine that have triggered message or content filter rules, the GUI may display content that did not actually trigger the filter action (along with content that triggered the filter action). The GUI display should be used as a guideline for locating content matches, but does not necessarily reflect an exact list of content matches. This occurs because the GUI uses less strict content matching logic than is used in the filters. This issue applies only to the highlighting in the message body. The table that lists the matched strings in each part of the message along with the associated filter rule is correct.

Figure 9-2 Matched Content Viewed in the Policy Quarantine



Description and Examples of Message Filter Actions

The following section describes the various message filter actions in use and their examples.

- [Skip Remaining Message Filters Action, page 9-59](#)
- [Drop Action, page 9-59](#)
- [Bounce Action, page 9-60](#)
- [Encrypt Action, page 9-60](#)
- [Notify and Notify-Copy Actions, page 9-61](#)
- [Blind Carbon Copy Actions, page 9-63](#)
- [Quarantine and Duplicate Actions, page 9-65](#)
- [Alter Recipient Action, page 9-66](#)
- [Alter Delivery Host Action, page 9-67](#)
- [Alter Source Host \(Virtual Gateway address\) Action, page 9-68](#)

- [Archive Action, page 9-68](#)
- [Strip Header Action, page 9-69](#)
- [Insert Header Action, page 9-69](#)
- [Edit Header Text Action, page 9-70](#)
- [Edit Body Text Action, page 9-70](#)
- [HTML Convert Action, page 9-71](#)
- [Bounce Profile Action, page 9-72](#)
- [Bypass Anti-Spam System Action, page 9-72](#)
- [Bypass Anti-Virus System Action, page 9-73](#)
- [Bypass File Reputation Filtering and File Analysis System Actions, page 9-73](#)
- [Bypass Outbreak Filter Scanning Action, page 9-73](#)
- [Add Message Tag Action, page 9-74](#)
- [Add Log Entry Action, page 9-74](#)
- [URL Reputation Actions, page 9-75](#)
- [URL Category Actions, page 9-76](#)
- [No Operation, page 9-78](#)

Skip Remaining Message Filters Action

The `skip-filters` action ensures that the message skips any further processing from message filters and continues through the email pipeline. The message that incurs the `skip-filters` action will be subject to anti-spam scanning and anti-virus scanning, if it is available on the appliance. The `skip-filters` action is the default final action for message filters.

The following filter notifies `customercare@example.com` and then immediately delivers any message addressed to `boss@admin`.

```
bossFilter:

    if(rcpt-to == 'boss@admin$')

    {

        notify('customercare@example.com');

        skip-filters();

    }
```

Drop Action

The `drop` action discards a message without any delivery. The message is not returned to the sender, not sent to the intended recipient, nor processed further in any way.

The following filter first notifies `george@whitehouse.gov` and then discards any message where the subject begins with `SPAM`.

```
spamFilter:

    if(subject == '^SPAM.*')

    {

        notify('george@whitehouse.gov');

        drop();

    }
```

Bounce Action

The `bounce` action sends the message back to the sender (Envelope Sender) without further processing. The following filter returns (bounces) any message from an email address that ends in `@yahoo\\.com`.

```
yahooFilter:

    if(mail-from == '@yahoo\\.com$')

    {

        bounce();

    }
```

Encrypt Action

The `encrypt` action uses the configured encryption profile to deliver encrypted messages to email recipients.

The following filter encrypts messages if they contain the term `[encrypt]` in the subject:

```
Encrypt_Filter:

    if ( subject == '\\[encrypt\\]' )

    {

        encrypt('My_Encryption_Profile');

    }
```



Note

You must have a Cisco Encryption Appliance in your network or a hosted key service configured to use this filter action. You must also have configured an encryption profile to use this filter action.

S/MIME Sign or Encrypt on Delivery Action

The `smime-gateway-deferred` action performs an S/MIME signing or encryption of the message using the specified sending profile during the delivery. This means that the message continues to the next stage of processing, and when all processing is complete, the message is signed or encrypted and delivered.

The following filter performs an S/MIME encryption on all the outgoing messages from a particular sender during the delivery:

```
smime-deferred:if(mail-from ==
"user@example.com"){smime-gateway-deferred("smime-encrypt");}
```

S/MIME Sign or Encrypt Action

The `smime-gateway` action performs an S/MIME signing or encryption using the specified sending profile and delivers the message, skipping any further processing.

The following filter performs an S/MIME signing on all the outgoing messages from a particular sender and delivers them immediately:

```
smime-deliver-now:if(mail-from == "user@example.com"){smime-gateway("smime-sign");}
```

Notify and Notify-Copy Actions

The `notify` and `notify-copy` actions send an email summary of the message to the specified email address. The `notify-copy` action also sends a copy of the original message, similar to the `bcc-scan` action. The notification summary contains:

- The contents of the Envelope Sender and Envelope Recipient (`MAIL FROM` and `RCPT TO`) directives from the mail transfer protocol conversation for the message.
- The message headers of the message.
- The name of the message filter that matched the message.

You can specify the recipient, subject line, from address, and notification template. the following filter selects messages with sizes larger than 4 megabytes, sends a notification email of each matching message to `admin@example.com`, and finally discards the message:

```
bigFilter:

if(body-size >= 4M)

{

    notify('admin@example.com');

    drop();

}
```

Or

```
bigFilterCopy:

if(body-size >= 4M)
```

```

{
    notify-copy('admin@example.com');

    drop();
}

```

The Envelope Recipient parameter may be any valid email address (for example, `admin@example.com` in the example above), or alternatively, may be the action variable `$EnvelopeRecipients` (see [Action Variables, page 9-56](#)), which specifies all Envelope Recipients of the message:

```

bigFilter:

if(body-size >= 4M)

{

    notify('$EnvelopeRecipients');

    drop();

}

```

The `notify` action also supports up to three additional, optional arguments that allow you to specify the subject header, the Envelope Sender, and a pre-defined text resource to use for the notification message. These parameters must appear in order, so a subject must be provided if the Envelope Sender is to be set or a notification template specified.

The subject parameter may contain action variables (see [Action Variables, page 9-56](#)) that will be replaced with data from the original message. By default, the subject is set to `Message Notification`.

The Envelope Sender parameter may be any valid email address, or alternatively, may be the action variable `$EnvelopeFrom`, which will set the return path of the message to the same as the original message

The notification template parameter is the name of an existing notification template. For more information, see [Notifications, page 9-85](#).

This example extends the previous one, but changes the subject to look like `[bigFilter] Message too large`, sets the return path to be the original sender, and uses the “`message.too.large`” template:

```

bigFilter:

if (body-size >= 4M)

{

    notify('admin@example.com', '[bigFilter] Message too large',

        '$EnvelopeFrom', 'message.too.large');

    drop();

}

```

You can also use the `$MatchedContent` action variable to notify senders or administrators that a content filter was triggered. The `$MatchedContent` action variable displays the content that triggered the filter. For example, the following filter sends a notification to an administrator if the email contains ABA account information.

```
ABA_filter:

if (body-contains ('*aba')){

  notify('admin@example.com', '[$MatchedContent]Account Information Displayed');

}
```

Related Topics

- [Notification Template, page 9-63](#)

Notification Template

You can use the Text Resources page or the `textconfig` CLI command to configure custom notification templates as text resources for use with the `notify()` and `notify-copy()` actions. If you do not create a custom notification template, a default template is used. The default template includes message headers, but the custom notification template does not include message headers by default. To include message headers in the custom notification, include the `$AllHeaders` action variable.

For more information, see the “Text Resources” chapter.

In the following example, when a large message triggers the filter shown below, an email is sent to the intended recipients explaining that the message was too large:

```
bigFilter:

if (body-size >= 4M)

{

  notify('$EnvelopeRecipients', '[$FilterName] Message too large',

        '$EnvelopeFrom', 'message.too.large');

  drop();

}
```

Blind Carbon Copy Actions

The `bcc` action sends an anonymous copy of the message to a specified recipient. This is sometimes referred to as message replication. Because no mention of the copy is made in the original message and the anonymous copy will never successfully bounce back to the recipient, the original sender and recipients of the message will not necessarily know that the copy was sent.

The following filter sends a blind carbon copy to `mom@home.org` for each message addressed to `sue` from `johnny`:

```
momFilter:

    if ((mail-from == '^johnny$') and (rcpt-to == '^sue$'))

    {

        bcc('mom@home.org');

    }
```

The `bcc` action also supports up to three additional, optional arguments that allow you to specify the subject header and Envelope Sender to use on the copied message, as well as an alt-mailhost. These parameters must appear in order, so a subject must be provided if the Envelope Sender is to be set.

The subject parameter may contain action variables (see [Action Variables, page 9-56](#)) that will be replaced with data from the original message. By default, this is set to the subject of the original message (the equivalent of `$Subject`).

The Envelope Sender parameter may be any valid email address, or alternatively, may be the action variable `$EnvelopeFrom`, which will set the return path of the message to the same as the original message.

This example expands the previous one by setting the subject to be `[Bcc] <original subject>`, and the return path set to `badbounce@home.org`:

```
momFilter:

    if ((mail-from == '^johnny$') and (rcpt-to == '^sue$'))

    {

        bcc('mom@home.org', '[Bcc] $Subject', 'badbounce@home.org');

    }
```

The alt-mailhost is the fourth parameter:

```
momFilterAltM:

    if ((mail-from == '^johnny$') and (rcpt-to == '^sue$'))

    {

        bcc('mom@home.org', '[Bcc] $Subject', '$EnvelopeFrom',
'momaltmailserver.example.com');

    }
```


**Caution**

The `Bcc()`, `notify()`, and `bounce()` filter actions can allow viruses through your network. The blind carbon copy filter action creates a new message which is a full copy of the original message. The `notify` filter action creates a new message that contains the headers of the original message. While it is rare, headers can contain viruses. The `bounce` filter action creates a new message which contains the first 10k of the original message. In all three cases, the new message will not be processed by anti-virus or anti-spam scanning.

To send to multiple hosts, you can call the `bcc()` action multiple times:

multiplealthosts:

```
if (recv-listener == "IncomingMail")
{
    insert-header('X-ORIGINAL-IP', '$remote_ip');

    bcc ('$EnvelopeRecipients', '$Subject', '$EnvelopeFrom', '10.2.3.4');

    bcc ('$EnvelopeRecipients', '$Subject', '$EnvelopeFrom', '10.2.3.5');

    bcc ('$EnvelopeRecipients', '$Subject', '$EnvelopeFrom', '10.2.3.6');
}
```

Related Topics

- [The `bcc-scan\(\)` Action, page 9-65](#)

The `bcc-scan()` Action

The `bcc-scan` action functions similarly to the `bcc` action, except that the message that is sent is treated as a brand new message and is therefore sent through the entire email pipeline.

momFilter:

```
if ((mail-from == '^johnny$') and (rcpt-to == '^sue$'))
{
    bcc-scan('mom@home.org');
}
```

Quarantine and Duplicate Actions

The `quarantine('quarantine_name')` action flags a message for inclusion into a queue called a quarantine. For more information about quarantines, see the “Quarantines” chapter. The `duplicate-quarantine('quarantine_name')` action immediately places a copy of the message into the specified quarantine and the original message continues through the email pipeline. The quarantine name is case sensitive.

When flagged for quarantine, the message continues through the rest of the email pipeline. When the message reaches the end of the pipeline, if the message has been flagged for one or more quarantines then it enters those queues. Otherwise, it is delivered. Note that if the message does not reach the end of the pipeline, it is not placed in a quarantine.

Accordingly, if a message filter contains a `quarantine()` action followed by a `bounce()` or `drop()` action, the message will not enter the quarantine, since the final action prevents the message from reaching the end of the pipeline. The same is true if a message filter includes a quarantine action, but the message is later dropped by anti-spam or anti-virus scanning, or a content filter. The `skip-filters()` action causes the message to skip any remaining message filters, but content filters may still apply. For example, if a message filter flags a message for quarantine and also includes the `skip-filters()` action, the message skips all remaining message filters and will be quarantined, unless another action in the email pipeline causes the message to be dropped.

In the following example, the message is sent to the Policy quarantine if the message contains any words within the dictionary named “secret_word.”

```
quarantine_codenames:

    if (dictionary-match ('secret_words'))
    {
        quarantine('Policy');
    }
```

In the following example, suppose a company has an official policy to drop all .mp3 file attachments. If an inbound message has a .mp3 attachment, the attachment is stripped and the remaining message (original body and remaining attachments) is sent to the original recipient. Another copy of the original message with all attachments will be quarantined (sent to the Policy quarantine). If it is necessary to receive the blocked attachment(s), the original recipient would then request that the message be released from the quarantine.

```
strip_all_mp3s:

    if (attachment-filename == '(?i)\\.mp3$') {
        duplicate-quarantine('Policy');
        drop-attachments-by-name '(?i)\\.mp3$';
    }
```

Alter Recipient Action

The `alt-rcpt-to` action changes all recipients of the message to the specified recipient upon delivery.

The following filter sends all messages with an Envelope Recipient address that contain `.freelist.com` and changes all recipients for the message to `system-lists@myhost.com`:

```
freelistFilter:

    if(rcpt-to == '\\.freelist\\.com$')

    {

        alt-rcpt-to('system-lists@myhost.com');

    }
```

Alter Delivery Host Action

The `alt-mailhost` action changes the IP address for all recipients of the selected message to the numeric IP address or hostname given.



Note

The `alt-mailhost` action prevents a message classified as spam by an anti-spam scanning engine from being quarantined. The `alt-mailhost` action overrides the quarantine action and sends it to the specified mail host.

The following filter redirects recipient addresses to the host `example.com` for all messages.

```
localRedirectFilter:

    if(true)

    {

        alt-mailhost('example.com');

    }
```

Thus, a message directed to `joe@anywhere.com` is delivered to the mailhost at `example.com` with the Envelope To address `joe@anywhere.com`. Note that any additional routing information specified by the `smtproutes` command still affects the routing of the message. (See [Routing Email for Local Domains, page 24-1.](#))



Note

The `alt-mailhost` action does not support specifying a port number. To do this, add an SMTP route instead.

The following filter redirects all messages to `192.168.12.5`:

```
local2Filter:

    if(true)

    {
```

```
        alt-mailhost('192.168.12.5');  
    }  
}
```

Alter Source Host (Virtual Gateway address) Action

The `alt-src-host` action changes the source host for the message to the source specified. The source host consists of the IP interface or group of IP interfaces that the messages should be delivered from. If a group of IP interfaces is selected, the system round-robins through all of the IP interfaces within the group as the source interface when delivering email. In essence, this allows multiple Virtual Gateway addresses to be created on a single Cisco Email Security appliance. For more information, see [Configuring Mail Gateways for all Hosted Domains Using Virtual Gateway™ Technology, page 24-59](#).

The IP interface may only be changed to an IP interface or interface group currently configured in the system. the following filter creates a Virtual Gateway using the outbound (delivery) IP interface `outbound2` for all messages received from a remote host with the IP address `1.2.3.4`.

```
externalFilter:  
  
    if(remote-ip == '1.2.3.4')  
  
    {  
  
        alt-src-host('outbound2');  
  
    }  
}
```

The following filter uses the IP interface group `Group1` for all messages received from a remote host with the IP address `1.2.3.4`.

```
groupFilter:  
  
    if(remote-ip == '1.2.3.4')  
  
    {  
  
        alt-src-host('Group1');  
  
    }  
}
```

Archive Action

The `archive` action saves a copy of the original message, including all message headers and recipients into an mbox-format file on the appliance. The action takes a parameter that is the name of the log file in which to save the message. The system automatically creates a log subscription with the specified filename when you create the filter, or you can also specify an existing filter log file. After the filter and the filter log file are created, the filter log options may then be edited with the `filters -> logconfig` subcommand.



Note

The `logconfig` command is a subcommand of `filters`. See [Using the CLI to Manage Message Filters, page 9-89](#) for a full description of how to use this subcommand.

The mbox format is a standard UNIX mailbox format, and there are many utilities available to make viewing the messages easier. Most UNIX systems allow you to type “`mail -f mbox.filename`” to view the files. The mbox format is in plain text, so you can use a simple text editor to view the contents of the messages.

In the following example, a copy of the message is saved to a log named `joesmith` if the Envelope Sender matches `joesmith@yourdomain.com`:

```
logJoeSmithFilter:

    if(mail-from == '^joesmith@yourdomain\\.com$')

    {

        archive('joesmith');

    }
```

Strip Header Action

The `strip-header` action examines the message for a particular header and removes those lines from the message before delivering it. When there are multiple headers, all instances of the header are removed (for example, the “Received:” header.)

In the following example, all messages have the header `X-DeleteMe` removed before transmission:

```
stripXDeleteMeFilter:

    if (true)

    {

        strip-header('X-DeleteMe');

    }
```

When working with headers, remember that the current value of the header includes changes made during processing (such as with filter actions that add, remove, or modify message headings). See [Message Header Rules and Evaluation, page 9-5](#) for more information.

Insert Header Action

The `insert-header` action inserts a new header into a message. AsyncOS does not verify the compliance to standards of the header you insert; you are responsible for ensuring that the resulting message complies with Internet standards for email.

The following example inserts a header named `X-Company` with the value set to `My Company Name` if the header is not already found in the message:

```
addXCompanyFilter:

    if (not header('X-Company'))

    {
```

```

        insert-header('X-Company', 'My Company Name');
    }

```

The `insert-header()` action allows the use of non-ASCII characters in the text of the header, while restricting the header name to be ASCII (to comply with standards). The transport encoding will be quoted-printable to maximize the readability.

**Note**

The `strip-headers` and `insert-header` actions can be used in combination to rewrite any message headers in the original message. In some case, it is valid to have multiple instances of the same header (for example, `Received:`) where in other cases, multiple instances of the same header could confuse a MUA (for example, multiple `Subject:` headers.)

When working with headers, remember that the current value of the header includes changes made during processing (such as with filter actions that add, remove, or modify message headings). See [Message Header Rules and Evaluation, page 9-5](#) for more information.

Edit Header Text Action

The `edit-header-text` action allows you to rewrite specified header text using the regular expression substitution function. The filter matches the regular expression within the header and replaces it with a regular expression you specify.

For example, an email contains the following subject header:

Subject: SCAN Marketing Messages

The following filter removes the “SCAN” text, and leaves the text, “Marketing Messages”, in the header:

```

Remove_SCAN: if true
{
    edit-header-text ('Subject', '^SCAN\s*', '');
}

```

After the filter processes the message, it returns the following header:

Subject: Marketing Messages

Edit Body Text Action

The `edit-body-text()` message filter is similar to the `Edit-Header-Text()` filter, but it operates across the body of the message instead of one of the headers.

The `edit-body-text()` message filter uses the following syntax where the first parameter is the regular expression to search for and the second parameter is the replacement text:

```

Example: if true {
    edit-body-text("parameter 1",

```

```
"parameter 2");

}
```

The `edit-body-text()` message filter only works on the message body parts. For more information about whether a given MIME part is considered a message “body” or a message “attachment”, see [Message Bodies vs. Message Attachments, page 9-5](#).

The following example shows a URL removed from a message and replaced with the text, ‘URL REMOVED’:

```
URL_Replaced: if true {

    edit-body-text("(?i)(?:https?|ftp)://[^\s\>]+", "URL REMOVED");

}
```

The following example shows a social security number removed from the body of a message and replaced with the text, “XXX-XX-XXXX”:

```
ssn: if true {

    edit-body-text("(?!000)(?:[0-6]\\d{2}|7(?:[0-6]\\d|7[012]))([0-9]?)(?!00)\\d\\d\\d\\d1(?:0000)\\d{4}",

    "XXX-XX-XXXX");

}
```

**Note**

You cannot use smart identifiers with the `edit-body-text()` filter at this time.

HTML Convert Action

While RFC 2822 defines a text format for email messages, there are extensions (such as MIME) to provide the transport of other content within an RFC 2822 message. AsyncOS can now use the `html-convert()` message filter to convert HTML to plain text using the following syntax:

```
Convert_HTML_Filter:

if (true)

{

    html-convert();

}
```

The Cisco message filters make a determination on whether a given MIME part is considered a message “body” or a message “attachment”. The `html-convert()` filter only works on the message body parts. For more information about message bodies and attachments, see [Message Bodies vs. Message Attachments, page 9-5](#).

Depending on the format, the `html-convert()` filter uses different methods to strip the HTML from within the documents.

If the message is plain text (text/plain), the message passes through the filter unchanged. If the message is a simple HTML message (text/html), all the HTML tags are stripped out of the message and the resulting body replaces the HTML message. The lines are not reformatted, and the HTML is not rendered in plain text. If the structure is MIME (with a multipart/alternative structure) and it contains both a text/plain part and text/html part with the same content, the filter removes the text/html part of the message and leaves the text/plain part of the message. For all other MIME types (such as multipart/mixed), all HTML body parts are stripped of their tags and reinserted into the message.

When encountered in a message filter, the `html-convert()` filter action only tags the message to be processed but does not immediately make a change to the message structure. The changes to the message only take effect after all processing is complete. This allows the other filter actions to process the original message body prior to modification.

Bounce Profile Action

The `bounce-profile` action assigns a previously-configured bounce profile to the message. (See [Directing Bounced Email, page 24-35](#).) If the message is undeliverable, the bounce options configured via the bounce profile are used. Using this feature overrides the bounce profile assigned to the message from the listener’s configuration (if one is assigned).

The following filter example assigns the bounce profile “fastbounce” to all email sent with the header `X-Bounce-Profile: fastbounce`:

```
fastbounce:

    if (header ('X-Bounce-Profile') == 'fastbounce') {

        bounce-profile ('fastbounce');

    }
```

Bypass Anti-Spam System Action

The `skip-spamcheck` action instructs the system to allow the message to bypass any content-based anti-spam filtering configured on the system. This action does nothing to the message if no content-based anti-spam filtering is configured, or if the message was never flagged to be scanned for spam in the first place.

The following example allows messages that have a high SenderBase Reputation Score to bypass the content-based anti-spam filtering feature:

```
whitelist_on_reputation:

    if (reputation > 7.5)

    {
```



```
        skip-spamcheck();  
    }  
}
```

Related Topics

- [How to Configure the Appliance to Scan Messages for Spam, page 13-2](#)
- [Protecting Appliance-Generated Messages From the Spam Filter, page 13-14](#)

Bypass Anti-Virus System Action

The `skip-viruscheck` action instructs the system to allow the message to bypass any virus protection system configured on the system. This action does nothing to the message if there is no anti-virus system configured, or if the message was never flagged to be scanned for viruses in the first place.

The following example specifies that messages received on the listener “private_listener” should bypass the anti-spam and the anti-virus systems.

```
internal_mail_is_safe:  
  
    if (recv-listener == 'private_listener')  
    {  
  
        skip-spamcheck();  
  
        skip-viruscheck();  
  
    }  
}
```

Bypass File Reputation Filtering and File Analysis System Actions

The `skip-ampcheck` action instructs the system to allow message to bypass File Reputation Filtering and File Analysis configured on the system. This action does nothing to the message if File Reputation Filtering and File Analysis is not configured, or if the message was never flagged to be scanned for File Reputation Filtering and File Analysis in the first place.

The following example specifies that messages with PDF attachments should bypass File Reputation Filtering and File Analysis.

```
skip_amp_scan:  
if (attachment-filetype == 'pdf')  
{  
    skip-ampcheck();  
}
```

Bypass Outbreak Filter Scanning Action

The `skip-vofcheck` action instructs the system to allow the message to bypass the Outbreak Filters scanning. This action does nothing to the message if Outbreak Filters scanning is not enabled.

The following example specifies that messages received on the listener “private_listener” should bypass Outbreak Filter scanning.

```
internal_mail_is_safe:

    if (recv-listener == 'private_listener') Outbreak Filters
    {
        skip-vofcheck();
    }
```

Add Message Tag Action

The `tag-message` action inserts a custom term into an outgoing message to use with RSA Email DLP policy filtering. You can configure a RSA Email DLP policy to limit scanning to messages with the message tag. The message tag is not visible to recipients. The tag name can contain any combination of characters from the set `[a-zA-Z0-9_-.]`.

For information on configuring a DLP policy to filter messages, see the “Data Loss Prevention” chapter.

The following example inserts a message tag into a message with “[Encrypt]” in the subject. You can then create a DLP policy that will encrypt messages with this message tag before delivering them if Cisco Email Encryption is available:

```
Tag_Message:

    if (subject == '^\\[Encrypt\\]')
    {
        tag-message('Encrypt-And-Deliver');
    }
```

Add Log Entry Action

The `log-entry` action inserts customized text into the Text Mail logs at the `INFO` level. The text can include action variables. You can use this action to insert useful text for debugging purposes and information on why a message filter performed a certain action. The log entry also appears in message tracking.

The following example inserts a log entry explaining that message was bounced because it possibly contained confidential company information:

```
CompanyConfidential:

    if (body-contains('Company Confidential'))
    {
        log-entry('Message may have contained confidential information.');
```

```

        bounce();
    }

```

URL Reputation Actions

Use the reputation score of URLs in messages to modify the URLs or their behavior. For important details and examples, see [Modifying URLs in Messages: Using URL Reputation and URL Category Actions in Filters](#), page 15-8 in Chapter 15, “URL Filtering.”

No rule is needed with these actions.

In URL Reputation actions:

- `msg_filter_name`: is the name of this message filter.
- `min_score` and `max_score` are the minimum and maximum scores in the range for which the action should apply. The applicable range includes the values that you specify.
Minimum and maximum scores must be between -10.0 and 10.0.
- To specify an action when the reputation service does not provide a score, use the corresponding "no-reputation" version of the action, as shown in the following subsections.
- `whitelist` is the name of a defined URL list (via the `urllistconfig` command.) Specifying a `whitelist` is optional.
- In place of `Preserve_signed`, enter 0 or 1:
 - 1 - Apply this action to unsigned messages only
 - 0 - Apply this action to all messages

If you do not specify a `preserve_signed` value, the action is applied to unsigned messages only.

Related Topics

- [Replace URL with Text, Based on URL Reputation](#), page 9-75
- [Defang URL, Based on URL Reputation](#), page 9-76
- [Redirect URL to Cisco Security Proxy, Based on URL Reputation](#), page 9-76

Replace URL with Text, Based on URL Reputation

To take action when the reputation service provides a score:

Use the `url-reputation-replace` action.

The syntax of a filter using the `url-reputation-replace` action is:

```

<msg_filter_name>:
if <condition>
{url-reputation-replace(<min_score>, <max_score>, '<replace_text>', '<whitelist>',
<Preserve_signed>);}

```

Where `replace_text` is the text with which to replace the URL.

To take action when the reputation service does not provide a score:

Use the `url-no-reputation-replace` action.

The syntax of a filter using the `url-no-reputation-replace` action is:

```
<msg_filter_name>:
if <condition>
{url-no-reputation-replace ('<replace_text>', '<whitelist>', <Preserve_signed>);}
```

Where `replace_text` is the text with which to replace the URL.

Defang URL, Based on URL Reputation

To take action when the reputation service provides a score:

Use the `url-reputation-defang` action.

The syntax of a filter using the `url-reputation-defang` action is:

```
<msg_filter_name>:
if <condition>
{url-reputation-defang (<min_score>, <max_score>, '<whitelist>', <Preserve_signed>);}
```

To take action when the reputation service does not provide a score:

Use the `url-no-reputation-defang` action.

The syntax of a filter using the `url-no-reputation-defang` action is:

```
<msg_filter_name>:
if <condition>
{url-no-reputation-defang ('<whitelist>', <Preserve_signed>);}
```

Redirect URL to Cisco Security Proxy, Based on URL Reputation

To take action when the reputation service provides a score:

Use the `url-reputation-proxy-redirect` action.

The syntax of a filter using the `url-reputation-proxy-redirect` action is:

```
<msg_filter_name>:
if <condition>
{url-reputation-proxy-redirect (<min_score>, <max_score>, '<whitelist>',
<Preserve_signed>);}
```

To take action when the reputation service does not provide a score:

Use the `url-no-reputation-proxy-redirect` action.

The syntax of a filter using the `url-no-reputation-proxy-redirect` action is:

```
<msg_filter_name>:
if <condition>
{url-no-reputation-proxy-redirect ('<whitelist>', <Preserve_signed>);}
```

URL Category Actions

Use the categories of URLs in messages to modify the URLs or their behavior. For important details, see [Modifying URLs in Messages: Using URL Reputation and URL Category Actions in Filters](#), page 15-8 in Chapter 15, “URL Filtering.”

No rule is needed with these actions.

In all URL Category actions:

- `msg_filter_name`: is the name of the message filter.
- `category-name` is the URL category. Separate multiple categories with commas. To obtain correct category names, look at a URL Category condition or action in a Content Filter. For descriptions and examples of the categories, see [About URL Categories, page 15-13](#).
- `url_white_list` is the name of a defined URL list (via the `urllistconfig` command.)
- `unsigned-only`: Enter 0 or 1.
 - 1 - Apply this action to unsigned messages only
 - 0 - Apply this action to all messages

Related Topics

- [Replace URL with Text, Based on URL Category, page 9-77](#)
- [Defang URL, Based on URL Category, page 9-77](#)
- [Redirect URL to Cisco Security Proxy, Based on URL Category, page 9-77](#)

Replace URL with Text, Based on URL Category

The syntax of a filter using the `url-category-replace` action is

```
<msg_filter_name>:
if <condition>
{
url-category-replace(['<category-name1>', '<category-name2>', ...,
'<category-name3>'], '<replacement-text>', '<url_white_list>', <unsigned-only>);
}
```

Where `replacement-text` is the text that you want to use to replace the URL.

Defang URL, Based on URL Category

The syntax of a filter using the `url-category-defang` action is:

```
<msg_filter_name>:
if <condition>
{
url-category-defang(['<category-name1>', '<category-name2>', ..., '<category-name3>'],
'<url_white_list>', <unsigned-only>);
}
```

Redirect URL to Cisco Security Proxy, Based on URL Category

The syntax of a filter using the `url-category-proxy-redirect` action is:

```
<msg_filter_name>:
if <condition>
{
```

```
url-category-proxy-redirect(['<category-name1>', '<category-name2>', ...,
'<category-name3>'], '<url_white_list>', <unsigned-only>);
}
```

No Operation

The No Operation action performs a no-op, or no operation. You can use this action in a message filter if you do not want to use any of the other actions such as Notify, Quarantine, or Drop. For example, to understand the behavior of a new message filter that you created, you can use the No Operation action. After the message filter is operational, you can monitor the behavior of the new message filter using the Message Filters report page, and fine-tune the filter to match your requirements.

The following example shows how to use No Operation action in a message filter.

```
new_filter_test: if header-repeats ('subject', X, 'incoming') {no-op();}
```

Attachment Scanning

AsyncOS can strip attachments from messages that are inconsistent with your corporate policies, while still retaining the ability to deliver the original message.

You can filter attachments based on their specific file type, fingerprint, or based on the content of the attachment. Using the fingerprint to determine the exact type of attachment prevents users from renaming a malicious attachment extension (for example, .exe) to a more commonly used extension (for example, .doc) in the hope that the renamed file would bypass attachment filters.

When you scan attachments for content, the Stellant attachment scanning engine extracts data from attachment files to search for the regular expression. It examines both data and metadata in the attachment file. If you scan an Excel or Word document, the attachment scanning engine can also detect the following types of embedded files: .exe, .dll, .bmp, .tiff, .pcx, .gif, .jpeg, .png, and Photoshop images.

Related Topics

- [Message Filters for Scanning Attachments, page 9-78](#)
- [Image Analysis, page 9-80](#)
- [Configuring the Image Analysis Scanning Engine, page 9-80](#)
- [Configuring the Message Filter to Perform Actions Based on Image Analysis Results, page 9-83](#)
- [Notifications, page 9-85](#)
- [Examples of Attachment Scanning Message Filters, page 9-86](#)

Message Filters for Scanning Attachments

The message filter actions described in [Table 9-8](#) are *non-final* actions. (Attachments are dropped and the message processing continues.)

The optional comment is text that is added to the message, much like a footer, and it can contain Message Filter Action Variables (see [Examples of Attachment Scanning Message Filters](#), page 9-86).

Table 9-8 Message Filter Actions for Attachment Filtering

Action	Syntax	Description
Drop Attachments by Name	drop-attachments-by-name (<regular expression>[, <optional comment>])	Drops all attachments on messages that have a filename that matches the given regular expression. Archive file attachments (zip, tar) will be dropped if they contain a file that matches. See Examples of Attachment Scanning Message Filters , page 9-86.
Drop Attachments by Type	drop-attachments-by-type (<MIME type>[, <optional comment>])	Drops all attachments on messages that have a MIME type, determined by either the given MIME type or the file extension. Archive file attachments (zip, tar) will be dropped if they contain a file that matches.
Drop Attachments by File Type	drop-attachments-by-filetype (<fingerprint name>[, <optional comment>])	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 Table 9-6 Attachment Groups , page 9-54.
Drop Attachments by MIME Type	drop-attachments-by-mimetype (<MIME type>[, <optional comment>])	Drops all attachments on messages that have a given MIME type. This action does not attempt to ascertain the MIME type by file extension and so it also does not examine the contents of archives.
Drop Attachments by Size	drop-attachments-by-size (<number>[, <optional comment>])	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.
Attachment Scanning	drop-attachments-where-contains (<regular expression>[, <optional comment>])	Drops all attachments on message that contain the regular expression. Archive files (zip, tar) will be dropped if any of the files they contain match the regular expression pattern.
Drop Attachments by Dictionary Matches	drop-attachments-where-dictionary-match (<dictionary name>)	This filter action strips attachments based on matches to dictionary terms. If the terms in the MIME parts considered to be an attachment match a dictionary term (and the user-defined threshold is met), the attachment is stripped from the email. See Examples of Attachment Scanning Message Filters , page 9-86.

Image Analysis

Some messages contain images that you may wish to scan for inappropriate content. You can use the image analysis engine to search for inappropriate content in email. Image analysis is not designed to supplement or replace your anti-virus and anti-spam scanning engines. Its purpose is to enforce acceptable use by identifying inappropriate content in email. Use the image analysis scanning engine to quarantine and analyze mail and to detect trends.

After you configure AsyncOS for image analysis, you can use image analysis filter rules to perform actions on suspect or inappropriate emails. Image scanning allows you to scan the following types of attached files: JPEG, BMP, PNG, TIFF, GIF, TGA, ICO, and PCX. The image analyzer uses algorithms that measure skin color, body size and curvature to determine the probability that the graphic contains inappropriate content. When you scan image attachments, Cisco fingerprinting determines the file type, and the image analyzer uses algorithms to analyze the image content. If the image is embedded in another file, the Stellent scanning engine extracts the file. The Stellent scanning engine can extract images from many file types, including Word, Excel, and PowerPoint documents. The image analysis verdict is computed on the message as a whole. If the message does not include any images, the message receives a score of "0" which maps to a "clean" verdict. Therefore, a message without any images will receive a "clean" verdict.



Note

Images cannot be extracted from PDF files.

Configuring the Image Analysis Scanning Engine

To enable image analysis from the GUI:

Procedure

Step 1 Go to Security Services > IronPort Image Analysis.

Step 2 Click Enable.

A success message displays, and the verdict settings display.

Figure 9-3 Cisco Image Analysis Overview
IronPort Image Analysis

IronPort Image Analysis Overview			
IronPort Image Analysis:	Enabled		
Image Analysis Sensitivity:	65		
Skip Images:	Enabled, 100 pixels		
Verdict Ranges:	CLEAN	SUSPECT	INAPPROPRIATE
	0 - 49	50 - 74	75 - 100
Edit Settings...			

The image analysis filter rule allows you to determine the actions to take based on the following verdicts:

- **Clean:** The image is free of inappropriate content. The image analysis verdict is computed on the message as a whole, so a message without any images will receive a "clean" verdict if scanned.
- **Suspect:** The image may contain inappropriate content.
- **Inappropriate:** The image contains inappropriate content.

These verdicts represent a numeric value assigned by the image analyzer algorithm to determine probability of inappropriate content.

The following values are recommended:

- Clean: 0 to 49
- Suspect: 50 to 74
- Inappropriate: 75 to 100

You can fine-tune image scanning by configuring the sensitivity setting, which helps reduce the number of false positives. For example, if you find that you are getting false positives, you can decrease the sensitivity setting. Or, conversely, if you find that the image scanning is missing inappropriate content, you may want to set the sensitivity higher. The sensitivity setting is a value between 0 (no sensitivity) and 100 (highly sensitive). The default sensitivity setting of 65 is recommended.

Related Topics

- [Tuning Image Analysis Settings, page 9-81](#)

Tuning Image Analysis Settings

Procedure

- Step 1** Go to **Security Services > IronPort Image Analysis**.
- Step 2** Click **Edit Settings**.

Figure 9-4 *Edit IronPort Image Analysis Settings*

Image Analysis Settings

☒ Enable IronPort Image Analysis

Image Analysis Sensitivity: Enter a value between 0 (least sensitive) and 100 (most sensitive). The recommended value is 65.

Skip Images: ☒ Skip image analysis for images smaller than pixels

Verdict Ranges		
CLEAN 0 to 49	SUSPECT 50 to 74	INAPPROPRIATE 75 to 100

Clean	Suspect	Inappropriate
The image is given a verdict of "Clean." The recommended range is 0-49.	The image is given a verdict of "Suspect". Use this verdict to create a rule in content filters to manage these messages. The recommended range is 50-74.	The image is given a verdict of "Inappropriate". Use this verdict to create a rule in content filters to manage these messages. The recommended range is 75-100.

- Step 3** Configure the settings for image analysis sensitivity. The default sensitivity setting of 65 is recommended.
- Step 4** Configure the settings for Clean, Suspect, and Inappropriate verdicts.
- When you configure the value ranges, ensure that you do not overlap values and that you use whole integers.
- Step 5** Optionally, configure AsyncOS to bypass scanning images that do not meet a minimum size requirement (recommended). By default, this setting is configured for 100 pixels. Scanning images that are smaller than 100 pixels can sometimes result in false positives.

You can also enable image analysis settings from the CLI via the `imageanalysisconfig` command:

```
test.com> imageanalysisconfig
```

```
IronPort Image Analysis: Enabled
```

```
Image Analysis Sensitivity: 65
```

```
Verdict Ranges: Clean (0-49), Suspect(50-74), Inappropriate (75+)
```

```
Skip small images with size less than 100 pixels (width or height)
```

```
Choose the operation you want to perform:
```

```
- SETUP - Configure IronPort Image Analysis.
```

```
[>] setup
```

```
IronPort Image Analysis: Enabled
```

```
Would you like to use IronPort Image Analysis? [Y]>
```

```
Define the image analysis sensitivity. Enter a value between 0 (least
```

```
sensitive) and 100 (most sensitive). As sensitivity increases, so does the false positive  
rate. The default setting of 65 is recommended.
```

```
[65]>
```

```
Define the range for a CLEAN verdict. Enter the upper bound of the CLEAN range by  
entering a value between 0 and 98. The default setting of 49 is recommended.
```

```
[49]>
```

```
Define the range for a SUSPECT verdict. Enter the upper bound of the SUSPECT range by  
entering a value between 50 and 99. The default setting of 74 is recommended.
```

```
[74]>
```

```
Would you like to skip scanning of images smaller than a specific size? [Y]>
```

Please enter minimum image size to scan in pixels, representing either height or width of a given image.

[100]>

Related Topics

- [Viewing the Verdict Score of a Particular Message, page 9-83](#)

Viewing the Verdict Score of a Particular Message

To see the verdict score for a particular message, you can view the mail logs. The mail logs display the image name or file name, the score for a particular message attachment. In addition, the log displays information about whether the images in a file were scannable or unscannable. Note that information in the log describes the result for each message attachment, rather than each image. For example, if the message had a zip attachment that contained a JPEG image, the log entry would contain the name of the zip file rather than the name of the JPEG. Also, if the zip file included multiple images then the log entry would include the maximum score of all the images. The unscannable notation indicates whether any of the images were unscannable.

The log does not contain information about how the scores translate to a particular verdict (clean, suspect or inappropriate). However, because you can use mail logs to track the delivery of specific messages, you can determine by the actions performed on the messages whether the mail contained inappropriate or suspect images.

For example, the following mail log shows attachments dropped by message filter rules as a result of Image Analysis scanning:

```
Thu Apr 3 08:17:56 2009 Debug: MID 154 IronPort Image Analysis: image 'Unscannable.jpg' is unscannable.
```

```
Thu Apr 3 08:17:56 2009 Info: MID 154 IronPort Image Analysis: attachment 'Unscannable.jpg' score 0 unscannable
```

```
Thu Apr 3 08:17:56 2009 Info: MID 6 rewritten to MID 7 by drop-attachments-where-image-verdict filter 'f-001'
```

```
Thu Apr 3 08:17:56 2009 Info: Message finished MID 6 done
```

Configuring the Message Filter to Perform Actions Based on Image Analysis Results

Once you enable image analysis, you must create a message filter to perform different actions for different message verdicts. For example, you may wish to deliver messages with a clean verdict, but quarantine messages that are determined to have inappropriate content.

**Note**

Cisco recommends you do not drop or bounce messages with inappropriate or suspect verdicts. Instead, send copies of violations to a quarantine for later review and better understanding of trend analysis.

The following filter shows messages tagged if the content is inappropriate or suspect:

```
image_analysis: if image-verdict == "inappropriate" {  
  
  strip-header("Subject");  
  
  insert-header("Subject", "[inappropriate image] $Subject");  
  
}  
  
  else {  
  
    if image-verdict == "suspect" {  
  
      strip-header("Subject");  
  
      insert-header("Subject", "[suspect image] $Subject");  
  
    }  
  
  }  
}
```

Related Topics

- [Creating Content Filters to Strip Attachments Based on Image Analysis Verdicts, page 9-84](#)

Creating Content Filters to Strip Attachments Based on Image Analysis Verdicts

After you enable image analysis, you can create a content filter to strip attachments based on image analysis verdicts, or you can configure a filter to perform different actions for different message verdicts. For example, you might decide to quarantine messages that contain inappropriate content.

To strip attachments based on image analysis verdicts:

Procedure

-
- Step 1** Click Mail Policies > Incoming Content Filters.
 - Step 2** Click Add Filter.
 - Step 3** Enter a name for the content filter.
 - Step 4** Under Actions, click **Add Action**.
 - Step 5** Under Strip Attachment by File Info, click **Image Analysis Verdict is:**
 - Step 6** Select from the following image analysis verdicts:
 - Suspect
 - Inappropriate
 - Suspect or Inappropriate
 - Unscannable

- Clean

To configure an action based on image analysis verdicts:

Procedure

-
- Step 1** Click Mail Policies > Incoming Content Filters.
- Step 2** Click Add Filter.
- Step 3** Enter a name for the content filter.
- Step 4** Under Conditions, click **Add Condition**.
- Step 5** Under Attachment File Info, click **Image Analysis Verdict**.
- Step 6** Choose from one of the following verdicts:
- Suspect
 - Inappropriate
 - Suspect or Inappropriate
 - Unscannable
 - Clean
- Step 7** Click **Add Action**.
- Step 8** Select an action to perform on messages based on the image analysis verdict.
- Step 9** Submit and commit your changes.
-

Notifications

Using the Text Resources page in the GUI or the `textconfig` CLI command to configure custom notification templates as text resources is another useful tool when used in conjunction with attachment filtering rules. The notification template supports non-ASCII characters (you are prompted to choose an encoding while creating the template).

In the following example, the `textconfig` command was first used to create a notification template named `strip.mp3` that will be inserted into to the body of the notification message. Then, an attachment filtering rule is created so that when an `.mp3` file has been stripped from a message, a notification email is sent to the intended recipients explaining that the `.mp3` file has been deleted.

```
drop-mp3s:

if (attachment-type == '*/mp3')

{ drop-attachments-by-filetype('Media');

    notify ('$EnvelopeRecipients', 'Your mp3 has been removed', '$EnvelopeFrom',
    'strip.mp3');

}
```

For more information, see [Notify and Notify-Copy Actions, page 9-61](#).

Examples of Attachment Scanning Message Filters

The following examples shows actions performed on attachments:

- [Inserting Headers, page 9-86](#)
- [Dropping Attachments by File Type, page 9-86](#)
- [Dropping Attachments by Dictionary Matches, page 9-88](#)
- [Quarantining Protected Attachments, page 9-88](#)
- [Detecting Unprotected Attachments, page 9-89](#)

Inserting Headers

In these examples, AsyncOS inserts headers when the attachments contain specified content.

In the following example, all of the attachments on the message are scanned for a keyword. If the keyword is present in all of the attachments, a custom X-Header is inserted:

```
attach_disclaim:

    if (every-attachment-contains('[dD]isclaimer') ) {

        insert-header("X-Example-Approval", "AttachOK");

    }
```

In the following example, the attachment is scanned for a pattern in the binary data. The filter uses the `attachment-binary-contains` filter rule to search for a pattern that indicates that the PDF document is encrypted. If the pattern is present in the binary data, a custom header is inserted:

```
match_PDF_Encrypt:

if (attachment-filetype == 'pdf' AND

attachment-binary-contains('/Encrypt')){

strip-header ('Subject');

insert-header ('Subject', '[Encrypted] $Subject');

}
```

Dropping Attachments by File Type

In the following example, the “executable” group of attachments (`.exe`, `.dll`, and `.scr`) is stripped from messages and text is added to the message, listing the filenames of the dropped files (via the `$dropped_filename` action variable). Note that the `drop-attachments-by-filetype` action examines

attachments and strips them based on the fingerprint of the file, and not just the three-letter filename extension. Note also that you can specify a single file type (“mpeg”) or you can refer to all of the members of the file type (“Media”):

```
strip_all_exes: if (true) {

    drop-attachments-by-filetype ('Executable', "Removed attachment:
$dropped_filename");

}
```

In the following example, the same “executable” group of attachments (.exe, .dll, and .scr) are stripped from messages whose Envelope Sender is not within the domain example.com.

```
strip_inbound_exes: if (mail-from != "@example\\.com$") {

    drop-attachments-by-filetype ('Executable');

}
```

In the following example, a specific member of a file type (“wmf”) as well as a the same “executable” group of attachments (.exe, .dll, and .scr) are stripped from messages whose Envelope Sender is not within the domain example.com.

```
strip_inbound_exes_and_wmf: if (mail-from != "@example\\.com$") {

    drop-attachments-by-filetype ('Executable');

    drop-attachments-by-filetype ('x-wmf');

}
```

In the following example, the “executable” pre-defined group of attachments is extended to include more attachment names. (Note that this action will *not* examine the attachments’ file type.)

```
strip_all_dangerous: if (true) {

    drop-attachments-by-filetype ('Executable');

    drop-attachments-by-name('( ?i)\\. (cmd|pif|bat)$');

}
```

The drop-attachments-by-name action supports non-ASCII characters.



Note

The drop-attachments-by-name action matches the regular expression against the filename captured from the MIME header. The filename captured from the MIME header may contain trailing spaces.

In the following example, a message is dropped if the attachment is not an .exe executable file type. However, the filter will not perform any action on the message if there is at least one attachment with the file type you want to filter out. For example, the following filter drops any message with an attachment that is not an .exe file type:

```
exe_check: if (attachment-filetype != "exe") {

    drop();

}
```

If a message has multiple attachments, the Email Security appliance does not drop the message if at least one of the attachments is an .exe file, even if the other attachments not .exe files.

Dropping Attachments by Dictionary Matches

This `drop-attachments-where-dictionary-match` action strips attachments based on matches to dictionary terms. If the terms in the MIME parts considered to be an attachment match a dictionary term (and the user-defined threshold is met), the attachment is stripped from the email. The following example shows attachment drops if words in the “secret_words” dictionary are detected in the attachment. Note that the threshold for the matches is set to one:

```
Data_Loss_Prevention: if (true) {

drop-attachments-where-dictionary-match("secret_words", 1);

}
```

Quarantining Protected Attachments

The `attachment-protected` filter tests whether any attachment in the message is password protected. You might use this filter on incoming mail to ensure that the attachments are scannable. According to this definition, a zip file containing one encrypted member along with unencrypted members will be considered protected. Similarly, PDF file that has no open password will not be considered protected, even though it may restrict copying or printing with a password. The following example shows protected attachments sent to a policy quarantine:

```
quarantine_protected:

if attachment-protected

{

quarantine("Policy");

}
```


Detecting Unprotected Attachments

The `attachment-unprotected` filter tests whether any attachment in the message is *not* password protected. This message filter complements the `attachment-protected` filter. You might use this filter on outgoing mail to detect outgoing mail that is unprotected. The following example shows AsyncOS detecting unprotected attachments on an outgoing listener and quarantining the messages:

```
quarantine_unprotected:

if attachment-unprotected

{

quarantine("Policy");

}
```

Using the CLI to Manage Message Filters

You can use the CLI to add, delete, activate and de-activate, import and export, and set logging options for message filters. The table below shows a summary of the commands and subcommands. The table below shows a summary of the commands and subcommands.

Table 9-9 **Message Filters Subcommands**

Syntax	Description
filters	The main command. This command is interactive; it asks you for more information (for example, <code>new</code> , <code>delete</code> , <code>import</code>).
new	Creates a new filter. If no location is given, it is appended to the current sequence. Otherwise, the filter will be inserted into the specific place in the sequence. For more information, see Creating a New Message Filter, page 9-90 .
delete	Deletes a filter by name or by sequence number. For more information, see Deleting a Message Filter, page 9-91 .
move	Rearranges the existing filters. For more information, see Moving a Message Filter, page 9-91 .
set	Sets filter to active or inactive state. For more information, see Activating and Deactivating a Message Filter, page 9-91 .
import	Replaces the current set of filters with a new set stored in a file (in the /configuration directory of the appliance). For more information, see Importing Message Filters, page 9-95 .
export	Exports the current set of filters to a file (in the /configuration directory of the appliance). For more information, see Exporting Message Filters, page 9-96 .
list	Lists information about a filter or filters. For more information, see Displaying a Message Filter List, page 9-96 .

Table 9-9 Message Filters Subcommands (continued)

Syntax	Description
detail	Prints detailed information about a specific filter, including the body of the filter rule itself. For more information, see Displaying Message Filter Details, page 9-96 .
logconfig	Enters the logconfig submenu of filters, allowing you to edit the log subscriptions from <code>archive()</code> filter actions. For more information, see Configuring Filter Log Subscriptions, page 9-96 .

**Note**

You must issue the `commit` command for filters to take effect.

Three types of parameters are:

Table 9-10 Filter Management Parameters

seqnum	An integer representing a filter based on its position in the list of filters. A <i>seqnum</i> of 2 represents the second filter in the list, for example.
filtname	The colloquial name of a filter.
range	A range may be used to represent more than one filter, and appears in the form of <i>X-Y</i> , where <i>X</i> and <i>Y</i> are the first and last <i>seqnums</i> that identify the extent. For example, <i>2-4</i> represents filters in the second, third, and fourth positions. Either <i>X</i> or <i>Y</i> may be left off to represent an open-ended list. For example, <i>-4</i> represents the first four filters, and <i>2-</i> represents all filters except the first. You can also use the keyword <code>all</code> to represent all the filters in the filter list.

Related Topics

- [Creating a New Message Filter, page 9-90](#)
- [Deleting a Message Filter, page 9-91](#)
- [Moving a Message Filter, page 9-91](#)
- [Activating and Deactivating a Message Filter, page 9-91](#)
- [Importing Message Filters, page 9-95](#)
- [Exporting Message Filters, page 9-96](#)
- [Viewing Non-ASCII Character Sets, page 9-96](#)
- [Displaying a Message Filter List, page 9-96](#)
- [Displaying Message Filter Details, page 9-96](#)
- [Configuring Filter Log Subscriptions, page 9-96](#)
- [Changing Message Encoding, page 9-98](#)
- [Sample Message Filters, page 9-100](#)

Creating a New Message Filter

```
new [seqnum|filtname|last]
```

Specifies the position at which to insert the new filter(s). If omitted, or given the keyword `last`, the filters entered in are appended to the list of filters. No gaps in the sequence numbers are allowed; you are not allowed to enter a *seqnum* outside the boundaries of the current list. If you enter an unknown *filename*, you are prompted to enter a valid *filename*, *seqnum*, or `last`.

After a filter has been entered, you may manually enter the filter script. When you are finished typing, end the entry by typing a period (.) on a line by itself.

The following conditions can cause errors:

- Sequence number beyond the current range of sequence numbers.
- Filter with a non-unique *filename*.
- Filter with a *filename* that is a reserved word.
- Filter with a syntax error.
- Filter with actions referring to non-existent system resources such as interfaces.

Deleting a Message Filter

```
delete [seqnum|filename|range]
```

Deletes the filter(s) identified.

The following conditions can cause errors:

- No filter with a given filter name.
- No filter with a given sequence number.

Moving a Message Filter

```
move [seqnum|filename|range seqnum|last]
```

Moves the filters identified by the first parameter to the position identified by the second parameter. If the second parameter is the keyword `last`, the filters are moved to the end of the list of filters. If more than one filter is being moved, their ordering remains the same in relation to one another.

The following conditions can cause errors:

- No filter with a given filter name.
- No filter with a given sequence number.
- Sequence number beyond the current range of sequence numbers.
- Movement would result in no change of sequence.

Activating and Deactivating a Message Filter

A given message filter is either *active* or *inactive* and it is also either *valid* or *invalid*. A message filter is only used for processing if it is both *active* and *valid*. You change an existing filter from active to inactive (and back again) via the CLI. A filter is invalid if it refers to a listener or interface which does not exist (or has been removed).

**Note**

You can determine if a filter is inactive by its syntax; AsyncOS changes the colon after the filter name to an exclamation point for inactive filters. If you use this syntax when entering or importing a filter, AsyncOS marks the filter as inactive.

For example, the following benign filter named “filterstatus” is entered. It is then made inactive using the `filter -> set` subcommand. Note that when the details of the filter are shown, the colon has been changed to an exclamation point (and is bold in the following example).

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

Choose the operation you want to perform:

- NEW - Create a new filter.
- IMPORT - Import a filter script from a file.

```
[> new
```

Enter filter script. Enter '.' on its own line to end.

```
filterstatus: if true{skip-filters();}
```

```
.
```

```
1 filters added.
```

Choose the operation you want to perform:

- NEW - Create a new filter.
- DELETE - Remove a filter.
- IMPORT - Import a filter script from a file.
- EXPORT - Export filters to a file
- MOVE - Move a filter to a different position.
- SET - Set a filter attribute.
- LIST - List the filters.
- DETAIL - Get detailed information on the filters.
- LOGCONFIG - Configure log subscriptions used by filters.
- ROLLOVERNOW - Roll over a filter log file.

```
[> list
```

```
Num Active Valid Name
1 Y Y filterstatus
```

Choose the operation you want to perform:

- NEW - Create a new filter.
- DELETE - Remove a filter.
- IMPORT - Import a filter script from a file.
- EXPORT - Export filters to a file
- MOVE - Move a filter to a different position.
- SET - Set a filter attribute.
- LIST - List the filters.
- DETAIL - Get detailed information on the filters.
- LOGCONFIG - Configure log subscriptions used by filters.
- ROLLOVERNOW - Roll over a filter log file.

```
[> set
```

Enter the filter name, number, or range:

```
[all]> all
```

Enter the attribute to set:

```
[active]> inactive
```

1 filters updated.

Choose the operation you want to perform:

- NEW - Create a new filter.
- DELETE - Remove a filter.
- IMPORT - Import a filter script from a file.

- EXPORT - Export filters to a file
- MOVE - Move a filter to a different position.
- SET - Set a filter attribute.
- LIST - List the filters.
- DETAIL - Get detailed information on the filters.
- LOGCONFIG - Configure log subscriptions used by filters.
- ROLLOVERNOW - Roll over a filter log file.

[> **detail**

Enter the filter name, number, or range:

[> **all**

Num Active Valid Name

1 N Y filterstatus

filterstatus! if (true) {

skip-filters();

}

Choose the operation you want to perform:

- NEW - Create a new filter.
- DELETE - Remove a filter.
- IMPORT - Import a filter script from a file.
- EXPORT - Export filters to a file
- MOVE - Move a filter to a different position.
- SET - Set a filter attribute.
- LIST - List the filters.
- DETAIL - Get detailed information on the filters.

```
- LOGCONFIG - Configure log subscriptions used by filters.  
- ROLLOVERNOW - Roll over a filter log file.  
  
[]>
```

Related Topics

- [Activating or Deactivating a Message Filter, page 9-95](#)

Activating or Deactivating a Message Filter

```
set [seqnum|filename|range] active|inactive
```

Sets the filters identified to have the given state. Legal states are:

- active: Set the state of the selected filters to be active.
- inactive: Set the state of the selected filters to be inactive.

The following conditions can cause errors:

- No filter with a given *filename*.
- No filter with a given sequence number.



Note

A filter which is inactive may also be noted in its syntax; the colon after the label (name of the filter) is changed to an exclamation point (!). A filter entered manually from the CLI, or imported, that contains this syntax, will automatically be marked inactive. For example, `mailfrompm!` instead of `mailfrompm:` is displayed.

Importing Message Filters

```
import filename
```

The name of the file containing filters to be processed. This file must reside in the configuration directory of the FTP/SCP root directory on the appliance, if you enabled FTP/SCP access for the interface with the `interfaceconfig` command. It is ingested and parsed, and any errors are reported. The filters imported replace all filters existing in the current filter set. See [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) for more information. Consider exporting the current filter list (see [Exporting Message Filters, page 9-96](#)) and then editing that file before importing.

When importing message filters, you are prompted to select the encoding used.

The following conditions can cause errors:

- File does not exist.
- Filter with a non-unique filter name.
- Filter with a *filename* that is a reserved word.
- Filter with a syntax error.
- Filter with actions referring to non-existent system resources such as interfaces.

Exporting Message Filters

```
export filename [seqnum|filename|range]
```

Output a formatted version of the existing filter set to a file in the configuration directory of the FTP/SCP root directory on the appliance. See [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) for more information.

When exporting message filters, you are prompted to select the encoding used.

The following conditions can cause errors:

- No filter with a given filter name.
- No filter with a given sequence number.

Viewing Non-ASCII Character Sets

The system displays filters containing non-ASCII characters in the CLI in UTF-8. If your terminal/display does not support UTF-8, the filter will be unreadable.

The best way to manage non-ASCII characters in filters is to edit the filter in a text file and then import that text file (see [Importing Message Filters, page 9-95](#)) into the appliance.

Displaying a Message Filter List

```
list [seqnum|filename|range]
```

Shows summarized information about the identified filters in a tabular form without printing the filter body. The information displayed includes:

- Filter name
- Filter sequence number
- Filter's active/inactive state
- Filter's valid/invalid state

The following conditions can cause errors:

- Illegal range format.

Displaying Message Filter Details

```
detail [seqnum|filename|range]
```

Provides full information about the identified filters, including the body of the filter and any additional state information.

Configuring Filter Log Subscriptions

```
logconfig
```


Enters a submenu that allows you to configure the filter log options for the mailbox files generated by the `archive()` action. These options are very similar to those used by the regular `logconfig` command, but the logs may only be created or deleted by adding or removing filters that reference them.

Each filter log subscription has the following default values, which can be modified using the `logconfig` subcommand:

- Retrieval method - FTP Poll
- File size - 10MB
- Max number of files - 10

For more information, see the “Logging” chapter.

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

Choose the operation you want to perform:

- NEW - Create a new filter.
- DELETE - Remove a filter.
- IMPORT - Import a filter script from a file.
- EXPORT - Export filters to a file
- MOVE - Move a filter to a different position.
- SET - Set a filter attribute.
- LIST - List the filters.
- DETAIL - Get detailed information on the filters.
- LOGCONFIG - Configure log subscriptions used by filters.
- ROLLOVERNOW - Roll over a filter log file.

```
[> logconfig
```

Currently configured logs:

```
1. "joesmith" Type: "Filter Logs" Retrieval: FTP Poll
```

Choose the operation you want to perform:

- EDIT - Modify a log setting.

```
[> edit
```

Enter the number of the log you wish to edit.

```
[> 1
```

Choose the method to retrieve the logs.

1. FTP Poll

2. FTP Push

3. SCP Push

```
[1]> 1
```

Please enter the filename for the log:

```
[joesmith.mbox]>
```

Please enter the maximum file size:

```
[10485760]>
```

Please enter the maximum number of files:

```
[10]>
```

Currently configured logs:

1. "joesmith" Type: "Filter Logs" Retrieval: FTP Poll

Enter "EDIT" to modify or press Enter to go back.

```
[>
```

Changing Message Encoding

You can use the `localeconfig` command to set the behavior of AsyncOS regarding modifying the encoding of message headings and footers during message processing:

```
example.com> localeconfig
```

Behavior when modifying headers: Use encoding of message body

Behavior for untagged non-ASCII headers: Impose encoding of message body

Behavior for mismatched footer or heading encoding: Only try encoding from message body

Choose the operation you want to perform:

- SETUP - Configure multi-lingual settings.

[> **setup**

If a header is modified, encode the new header in the same encoding as

the message body? (Some MUAs incorrectly handle headers encoded in a

different encoding than the body. However, encoding a modified header

in the same encoding as the message body may cause certain characters in the modified header to be lost.) [Y]>

If a non-ASCII header is not properly tagged with a character set and

is being used or modified, impose the encoding of the body on the

header during processing and final representation of the message?

(Many MUAs create non-RFC-compliant headers that are then handled in

an undefined way. Some MUAs handle headers encoded in character sets

that differ from that of the main body in an incorrect way. Imposing the encoding of the body on the header may encode

the header more precisely. This will be used to interpret the content of headers for processing, it will not modify or rewrite the header

unless that is done explicitly as part of the processing.) [Y]>

Footers or headings are added in-line with the message body whenever

possible. However, if the footer or heading is encoded differently

than the message body, and if imposing a single encoding will cause

loss of characters, it will be added as an attachment. The system will

```

always try to use the message body's encoding for the footer or
heading. If that fails, and if the message body's encoding is US-
ASCII, the system can try to edit the message body to use the footer's
or heading's encoding. Should the system try to impose the footer's
or headings's encoding on the message body? [N]> y

```

```

Behavior when modifying headers: Use encoding of message body
Behavior for untagged non-ASCII headers: Impose encoding of message
body. Behavior for mismatched footer or heading encoding: Try both
body and footer or heading encodings

```

```

Choose the operation you want to perform:
- SETUP - Configure multi-lingual settings.

```

The first prompt determines whether or not a message header's encoding should be changed to match that of the message body if the header is changed (via a filter, for example).

The second prompt controls whether or not the appliance should impose the encoding of the message body on the header if the header is not properly tagged with a character set.

The third prompt is used to configure how disclaimer stamping (and multiple encodings) in the message body works. Please see “Disclaimer Stamping and Multiple Encodings” in the “Text Resources” chapter for more information.

Sample Message Filters

In the following example, the `filter` command is used to create three new filters:

- The first filter is named **big_messages**. It uses the `body-size` rule to drop messages larger than 10 megabytes.
- The second filter is named **no_mp3s**. It uses the `attachment-filename` rule to drop messages that contain attachments with the filename extension of `.mp3`.
- The third filter is named **mailfrompm**. It uses `mail-from` rule examines all mail from `postmaster@example.com` and blind-carbon copies `administrator@example.com`.

Using the `filter -> list` subcommand, the filters are listed to confirm that they are active and valid, and then the first and last filters are switched in position using the `move` subcommand. Finally, the changes are committed so that the filters take effect.

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

Choose the operation you want to perform:

- NEW - Create a new filter.
- IMPORT - Import a filter script from a file.

```
[> new
```

Enter filter script. Enter '.' on its own line to end.

```
big_messages:
```

```
    if (body-size >= 10M) {  
        drop();  
    }
```

```
.
```

```
1 filters added.
```

Choose the operation you want to perform:

- NEW - Create a new filter.
- DELETE - Remove a filter.
- IMPORT - Import a filter script from a file.
- EXPORT - Export filters to a file
- MOVE - Move a filter to a different position.
- SET - Set a filter attribute.
- LIST - List the filters.
- DETAIL - Get detailed information on the filters.
- LOGCONFIG - Configure log subscriptions used by filters.
- ROLLOVERNOW - Roll over a filter log file.

```
[> new
```

Enter filter script. Enter '.' on its own line to end.

no_mp3s:

```
if (attachment-filename == '(?i)\\.mp3$') {
    drop();
}
```

.

1 filters added.

Choose the operation you want to perform:

- NEW - Create a new filter.
- DELETE - Remove a filter.
- IMPORT - Import a filter script from a file.
- EXPORT - Export filters to a file
- MOVE - Move a filter to a different position.
- SET - Set a filter attribute.
- LIST - List the filters.
- DETAIL - Get detailed information on the filters.
- LOGCONFIG - Configure log subscriptions used by filters.
- ROLLOVERNOW - Roll over a filter log file.

[> **new**

Enter filter script. Enter '.' on its own line to end.

mailfrompm:

```
if (mail-from == "^postmaster$")
{ bcc ("administrator@example.com");}
```

.

1 filters added.

Choose the operation you want to perform:

- NEW - Create a new filter.
- DELETE - Remove a filter.
- IMPORT - Import a filter script from a file.
- EXPORT - Export filters to a file
- MOVE - Move a filter to a different position.
- SET - Set a filter attribute.
- LIST - List the filters.
- DETAIL - Get detailed information on the filters.
- LOGCONFIG - Configure log subscriptions used by filters.
- ROLLOVERNOW - Roll over a filter log file.

[> **list**

Num Active Valid Name

1	Y	Y	big_messages
2	Y	Y	no_mp3s
3	Y	Y	mailfrompm

Choose the operation you want to perform:

- NEW - Create a new filter.
- DELETE - Remove a filter.
- IMPORT - Import a filter script from a file.
- EXPORT - Export filters to a file
- MOVE - Move a filter to a different position.
- SET - Set a filter attribute.
- LIST - List the filters.
- DETAIL - Get detailed information on the filters.
- LOGCONFIG - Configure log subscriptions used by filters.
- ROLLOVERNOW - Roll over a filter log file.

```
[> move
```

Enter the filter name, number, or range to move:

```
[> 1
```

Enter the target filter position number or name:

```
[> last
```

1 filters moved.

Choose the operation you want to perform:

- NEW - Create a new filter.
- DELETE - Remove a filter.
- IMPORT - Import a filter script from a file.
- EXPORT - Export filters to a file
- MOVE - Move a filter to a different position.
- SET - Set a filter attribute.
- LIST - List the filters.
- DETAIL - Get detailed information on the filters.
- LOGCONFIG - Configure log subscriptions used by filters.
- ROLLOVERNOW - Roll over a filter log file.

```
[> list
```

Num Active Valid Name

```
1   Y      Y   no_mp3s
2   Y      Y   mailfrompm
3   Y      Y   big_messages
```

Choose the operation you want to perform:

- NEW - Create a new filter.

- DELETE - Remove a filter.
- IMPORT - Import a filter script from a file.
- EXPORT - Export filters to a file
- MOVE - Move a filter to a different position.
- SET - Set a filter attribute.
- LIST - List the filters.
- DETAIL - Get detailed information on the filters.
- LOGCONFIG - Configure log subscriptions used by filters.
- ROLLOVERNOW - Roll over a filter log file.

```
[> move
```

Enter the filter name, number, or range to move:

```
[> 2
```

Enter the target filter position number or name:

```
[> 1
```

1 filters moved.

Choose the operation you want to perform:

- NEW - Create a new filter.
- DELETE - Remove a filter.
- IMPORT - Import a filter script from a file.
- EXPORT - Export filters to a file
- MOVE - Move a filter to a different position.
- SET - Set a filter attribute.
- LIST - List the filters.
- DETAIL - Get detailed information on the filters.
- LOGCONFIG - Configure log subscriptions used by filters.

- ROLLOVERNOW - Roll over a filter log file.

[> **list**

Num Active Valid Name

1	Y	Y	mailfrompm
2	Y	Y	no_mp3s
3	Y	Y	big_messages

Choose the operation you want to perform:

- NEW - Create a new filter.

- DELETE - Remove a filter.

- IMPORT - Import a filter script from a file.

- EXPORT - Export filters to a file

- MOVE - Move a filter to a different position.

- SET - Set a filter attribute.

- LIST - List the filters.

- DETAIL - Get detailed information on the filters.

- LOGCONFIG - Configure log subscriptions used by filters.

- ROLLOVERNOW - Roll over a filter log file.

[>

mail3.example.com> **commit**

Please enter some comments describing your changes:

[> **entered and enabled 3 filters: no_mp3s, mailfrompm, big_messages**

Do you want to save the current configuration for rollback? [Y]> n

Changes committed: Fri May 23 11:42:12 2014 GMT

Message Filter Examples

This section contains some real world examples of filters with a brief discussion of each.

Related Topics

- [Open-Relay Prevention Filter, page 9-107](#)
- [Policy Enforcement Filters, page 9-107](#)
- [Routing and Domain Spoofing, page 9-111](#)

Open-Relay Prevention Filter

This filter bounces messages with addresses using %, extra @, and ! characters in email addresses:

- `user%otherdomain@validdomain`
- `user@otherdomain@validdomain:`
- `domain!user@validdomain`

```
sourceRouted:
```

```
if (rcpt-to == "(%|@|!)(.*)@") {  
  
    bounce();  
  
}
```

Cisco appliances are not susceptible to these third party relay hacks that are often used to exploit traditional Sendmail/Qmail systems. As many of these symbols (for example %) can be part of a perfectly legal email address, Cisco appliances will accept these as valid addresses, verify them against the configured recipient lists, and pass them on to the next internal server. Cisco appliances do not relay these messages to the world.

These filters are put in place to protect users who may have open-source MTAs that are misconfigured to allow relay of these types of messages.



Note

You can also configure a listener to handle these types of addresses. See [Listening for Connection Requests by Creating a Listener via the GUI, page 5-8](#) for more information.

Policy Enforcement Filters

- [Notify Based on Subject Filter, page 9-108](#)
- [BCC and Scan Mail Sent to Competitors, page 9-108](#)
- [Block Specific User Filter, page 9-108](#)
- [Archive and Drop Messages Filter, page 9-109](#)
- [Large “To:” Header Filter, page 9-109](#)
- [Blank “From:” Filter, page 9-109](#)

- [SRBS Filter, page 9-110](#)
- [Alter SRBS Filter, page 9-110](#)
- [Filename Regex Filter, page 9-110](#)
- [Show SenderBase Reputation Score in Header Filter, page 9-111](#)
- [Insert Policy into Header Filter, page 9-111](#)
- [Too Many Recipients Bounce Filter, page 9-111](#)

Notify Based on Subject Filter

This filter sends notification based on whether the subject contains specific words:

```
search_for_sensitive_content:

if (Subject == "(?i)plaintiff|lawsuit|judge" ) {

    notify ("admin@company.com");

}
```

BCC and Scan Mail Sent to Competitors

This filter scans and blind copies messages that are sent to competitors. Note that you could use a dictionary and the header-dictionary-match() rule to specify a more flexible list of competitors (see [Dictionary Rules, page 9-36](#)):

```
competitorFilter:

if (rcpt-to == '@competitor1.com|@competitor2.com') {

    bcc-scan('legal@example.com');

}
```

Block Specific User Filter

Use this filter to block email from a specific address:

```
block_harrasing_user:

if (mail-from == "ex-employee@hotmail\\.com") {

    notify ("admin@company.com");

    drop ();

}
```

Archive and Drop Messages Filter

Log and drop only the messages that have matching filetypes:

```
drop_attachments:

if (mail-from != "user@example.com") AND (attachment-filename ==

'(?i)\.(asp|bas|bat|cmd|cpl|exe|hta|ins|isp|js)$')

{

    archive("Drop_Attachments");

    insert-header("X-Filter", "Dropped by: $FilterName MID: $MID");
drop-attachments-by-name("\.(asp|bas|bat|cmd|cpl|exe|hta|ins|isp|js)$");

}
```

Large “To:” Header Filter

Find messages with very large “To” headers.

Use the `archive()` line for verification of proper action, with `drop()` enabled or disabled for extra safety:

```
toTooBig:

if(header('To') == "^.{500,}") {

    archive('tooTooBigdropped');

    drop();

}
```

Blank “From:” Filter

Identify blank “From” headers,

This filter can alleviate various forms of blank “from” addresses:

```
blank_mail_from_stop:

if (recv-listener == "InboundMail" AND header("From") == "^$|<\\s*>") {

    drop ();

}
```

If you also want to drop messages with a blank envelope from, use this filter:

```
blank_mail_from_stop:

if (recv-listener == "InboundMail" AND (mail-from == "^$|<\\s*>" OR header ("From") ==
"^$|<\\s*>"))

{

    drop ();

}
```

SRBS Filter

SenderBase Reputation filter:

```
note_bad_reps:

if (reputation < -2) {

    strip-header ('Subject');

    insert-header ('Subject', '***BadRep $Reputation *** $Subject');

}
```

Alter SRBS Filter

Alter the (SenderBase Reputation Score) SBRS threshold for certain domains:

```
mod_sbrs:

if ( (rcpt-count == 1) AND (rcpt-to == "@domain\\.com$") AND (reputation < -2) ) {

    drop ();

}
```

Filename Regex Filter

This filter specifies a range of size for the body of the message, and looks for an attachment that matches the regular expression (this matches files named “readme.zip”, “readme.exe”, “attach.exe”, and so forth.):

```
filename_filter:

if ((body-size >= 9k) AND (body-size <= 20k)) {

    if (body-contains "(?i)(readme|attach|information)\\. (zip|exe)$") {

        drop ();

    }

}
```

```
    }
}
```

Show SenderBase Reputation Score in Header Filter

Remember to log the headers (see the “Logging” chapter) so they appear in the mail log:

```
Check_SBRs:

if (true) {

    insert-header('X-SBRs', '$Reputation');

}
```

Insert Policy into Header Filter

Show which mail flow policy accepted the connection:

```
Policy_Tracker:

if (true) {

    insert-header ('X-HAT', 'Sender Group $Group, Policy $Policy applied.');
```

```
}
```

Too Many Recipients Bounce Filter

Bounce all outbound email messages with more than 50 recipients from more than two unique domains:

```
bounce_high_rcpt_count:

if ( ( rcpt-count > 49) AND (rcpt-to != "@example\\.com$") ) {

    bounce-profile ("too_many_rcpt_bounce"); bounce ();

}
```

Routing and Domain Spoofing

- [Using Virtual Gateways Filter, page 9-112](#)
- [Same Listener for Deliver and Listener Filter, page 9-112](#)
- [Single Listener Filter, page 9-112](#)
- [Drop Spoofed Domain Filter \(Single Listener\), page 9-113](#)
- [Drop Spoofed Domain Filter \(Multiple Listeners\), page 9-113](#)
- [Another Drop Spoofed Domain Filter, page 9-113](#)

- [Detect Looping Filter, page 9-114](#)

Using Virtual Gateways Filter

Segment traffic using virtual gateways. Assuming you have two Interfaces on the system, 'public1' and 'public2', and the default delivery interface is 'public1'. This would force all of your outbound traffic over the second interface; since bounces and other similar types of mail do not go through filters, they will be delivered from public1:

```
virtual_gateways:

if (recv-listener == "OutboundMail") {

    alt-src-host ("public2");

}
```

Same Listener for Deliver and Listener Filter

Use the same listener for delivery and receiving. This filter will allow you to send any messages received on the public listener "listener1" out the interface "listener1" (you will have to set up a unique filter for each public listener configured):

```
same_listener:

if (recv-inj == 'listener1') {

    alt-src-host('listener1');

}
```

Single Listener Filter

Make the filter work on a single listener. For example, specify a specific listener for message filter processing instead of being performed system wide.

```
textfilter-new:

if (recv-inj == 'inbound' and body-contains("some spammy message")) {

    alt-rcpt-to ("spam.quarantine@spam.example.com");

}
```


Drop Spoofed Domain Filter (Single Listener)

Drop email with a spoofed domain (pretending to be from an internal address; works with a single listener). IP addresses below represent fictional domain for `mycompany.com`:

```
DomainSpoofed:

if (mail-from == "mycompany\\.com$") {

    if ((remote-ip != "1.2.") AND (remote-ip != "3.4.")) {

        drop();

    }

}
```

Drop Spoofed Domain Filter (Multiple Listeners)

As above, but works with multiple listeners:

```
domain_spoof:

if ((recv-listener == "Inbound") and (mail-from == "@mycompany\\.com")) {

    archive('domain_spoof');

    drop ();

}
```

Another Drop Spoofed Domain Filter

Summary: Anti domain spoof filter:

```
reject_domain_spoof:

if (recv-listener == "MailListener") {

    insert-header("X-Group", "$Group");

    if ((mail-from == "@test\\.mycompany\\.com") AND (header("X-Group") != "RELAYLIST")) {

        notify("me@here.com");

        drop();

        strip-header("X-Group");

    }

}
```

Detect Looping Filter

This filter is used to detect, stop, and determine what is causing, a mail loop. This filter can help determine a configuration issue on the Exchange server or elsewhere.

```
External_Loop_Count:

if (header("X-ExtLoop1")) {

    if (header("X-ExtLoopCount2")) {

        if (header("X-ExtLoopCount3")) {

            if (header("X-ExtLoopCount4")) {

                if (header("X-ExtLoopCount5")) {

                    if (header("X-ExtLoopCount6")) {

                        if (header("X-ExtLoopCount7")) {

                            if (header("X-ExtLoopCount8")) {

                                if (header("X-ExtLoopCount9")) {

                                    notify ('joe@example.com');

                                    drop();

                                }

                                else {insert-header("X-ExtLoopCount9", "from
                                    $RemoteIP");}}

                                else {insert-header("X-ExtLoopCount8", "from $RemoteIP");}}

                                else {insert-header("X-ExtLoopCount7", "from $RemoteIP");}}

                                else {insert-header("X-ExtLoopCount6", "from $RemoteIP");}}

                                else {insert-header("X-ExtLoopCount5", "from $RemoteIP");}}

                                else {insert-header("X-ExtLoopCount4", "from $RemoteIP");}}

                                else {insert-header("X-ExtLoopCount3", "from $RemoteIP");}}

                            else {insert-header("X-ExtLoopCount2", "from $RemoteIP");}}

                        else {insert-header("X-ExtLoop1", "1"); }
```



Note

By default, AsyncOS automatically detects mail loops and will drop messages after 100 loops.

Configuring Scan Behavior

You can control the behavior of body and attachment scanning, such as the attachment types to be skipped during a scan by configuring the scanning parameters. Use the Scan Behavior page or the `scanconfig` command to configure these parameters. Scan behavior settings are global settings, meaning that they affect the behavior of all the scans.



Note

If you want to scan a MIME type that may be included in a zip or compressed file, you must include list 'compressed' or 'zip' or 'application/zip' in the scan list.

Procedure

Step 1 Click **Security Services > Scan Behavior**.

Step 2 Define the attachment type mapping. Do one of the following:

- Add a new attachment type mapping. Click **Add Mapping**.
- Import a list of attachment type mappings using a configuration file. Click **Import List**, and import the desired configuration file from the configuration directory.



Note

To perform this step, the configuration file must be present in the configuration directory of your appliance. See [Managing the Configuration File, page 33-7](#).

- Click **Edit** to modify an existing attachment type mapping.

Step 3 Configure the global settings. Do the following:

- Under Global Settings, click **Edit Global Settings**.
- Edit the desired fields:

Table 9-11

Field	Description
Action for attachments with MIME types / fingerprints in table above	Choose whether to scan or skip attachments types defined in the attachment type mapping.
Maximum depth of attachment recursion to scan	Specify the level up to which the recursive attachments are to be scanned.
Maximum attachment size to scan	Specify the maximum size of attachments to scan.
Attachment Metadata scan	Specify whether to scan or skip metadata of the attachments.
Attachment scanning timeout	Specify the scanning time-out period.
Assume attachment matches pattern if not scanned for any reason	Specify whether to consider unscanned attachments as match to the search pattern.
Action when message cannot be deconstructed to remove specified attachments	Specify the action to be taken when a message could not be deconstructed to remove specified attachments.

Table 9-11

Field	Description
Bypass all filters in case of a content or message filter error	Specify whether to bypass all filters in case of a content or message filter error.
Encoding to use when none is specified	Specify the encoding to be used if no encoding is specified.
Convert opaque-signed messages to clear-signed (S/MIME unpacking)	Specify whether to convert opaque-signed messages to clear-signed (S/MIME unpacking).

c. Click **Submit**.

Step 4 Click **Commit Changes**.

Related Topics

- [Modifying Scan Parameters using CLI, page 9-116](#)

Modifying Scan Parameters using CLI

You can use the `scanconfig` command to configure the behavior of body and attachment scanning. In the following example, the `scanconfig` command sets the following parameters:

- MIME types of `video/*`, `audio/*`, `image/*` are not scanned for content.
- Nested (recursive) archive attachments up to 50 levels are scanned.
- The maximum size for attachments to be scanned is 25 megabytes; anything larger will be skipped. (The default is 5 megabytes.)
- The attachment is enabled for metadata scanning. When the scanning engine scans attachments, it scans the metadata for the regular expression. This is the default setting.
- The attachment timeout scanning is configured for 60 seconds. The default is 30 seconds.
- Attachments that were not scanned are assumed to not match the search pattern. (This is the default behavior.)
- The `application/(x-)pkcs7-mime` (opaque-signed) parts of a message are converted to `multipart/signed` (clear-signed) to provide the message's content for processing. The default is not to convert opaque-signed messages.



Note

When setting the `assume the attachment matches the search pattern` to `Y`, messages that cannot be scanned will cause the message filter rule to evaluate to true. This could result in unexpected behavior, such as the quarantining of messages that do not match a dictionary, but were quarantined because their content could not be correctly scanned.

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

```
There are currently 5 attachment type mappings configured to be SKIPPED.
```

Choose the operation you want to perform:

- NEW - Add a new entry.
- DELETE - Remove an entry.
- SETUP - Configure scanning behavior.
- IMPORT - Load mappings from a file.
- EXPORT - Save mappings to a file.
- PRINT - Display the list.
- CLEAR - Remove all entries.
- SMIME - Configure S/MIME unpacking.

[1]> **setup**

1. Scan only attachments with MIME types or fingerprints in the list.
2. Skip attachments with MIME types or fingerprints in the list.

Choose one:

[2]> **2**

Enter the maximum depth of attachment recursion to scan:

[5]> **10**

Enter the maximum size of attachment to scan:

[5242880]> **10m**

Do you want to scan attachment metadata? [Y]> **Y**

Enter the attachment scanning timeout (in seconds):

[30]> **60**

If a message has attachments that were not scanned for any reason (e.g. because of size, depth limits, or scanning timeout), assume the attachment matches the search pattern?
[N]>

If a message could not be deconstructed into its component parts in order to remove specified attachments, the system should:

1. Deliver
2. Bounce
3. Drop

[1]> **1**

Configure encoding to use when none is specified for plain body text or anything with MIME type plain/text or plain/html.

1. US-ASCII
2. Unicode (UTF-8)
3. Unicode (UTF-16)
4. Western European/Latin-1 (ISO 8859-1)
5. Western European/Latin-1 (Windows CP1252)
6. Traditional Chinese (Big 5)
7. Simplified Chinese (GB 2312)
8. Simplified Chinese (HZ GB 2312)
9. Korean (ISO 2022-KR)
10. Korean (KS-C-5601/EUC-KR)
11. Japanese (Shift-JIS (X0123))
12. Japanese (ISO-2022-JP)
13. Japanese (EUC)

[1]>

Scan behavior changed.

There are currently 5 attachment type mappings configured to be SKIPPED.

Choose the operation you want to perform:

- NEW - Add a new entry.
- DELETE - Remove an entry.
- SETUP - Configure scanning behavior.
- IMPORT - Load mappings from a file.
- EXPORT - Save mappings to a file.
- PRINT - Display the list.
- CLEAR - Remove all entries.
- SMIME - Configure S/MIME unpacking.

[> **SMIME**

Do you want to convert opaque-signed messages to clear-signed? This will provide the clear text content for various blades to process. [N]> Y

There are currently 5 attachment type mappings configured to be SKIPPED.

Choose the operation you want to perform:

- NEW - Add a new entry.
- DELETE - Remove an entry.
- SETUP - Configure scanning behavior.
- IMPORT - Load mappings from a file.
- EXPORT - Save mappings to a file.
- PRINT - Display the list.
- CLEAR - Remove all entries.
- SMIME - Configure S/MIME unpacking.

[> **print**

1. Fingerprint Image
2. Fingerprint Media
3. MIME Type audio/*
4. MIME Type image/*
5. MIME Type video/*

There are currently 5 attachment type mappings configured to be SKIPPED.

Choose the operation you want to perform:

- NEW - Add a new entry.
- DELETE - Remove an entry.
- SETUP - Configure scanning behavior.
- IMPORT - Load mappings from a file.
- EXPORT - Save mappings to a file.
- PRINT - Display the list.
- CLEAR - Remove all entries.
- SMIME - Configure S/MIME unpacking.

[]>



Mail Policies

- [Overview of Mail Policies, page 10-1](#)
- [How to Enforce Mail Policies on a Per-User Basis, page 10-2](#)
- [Handling Incoming and Outgoing Messages Differently, page 10-3](#)
- [Matching Users to a Mail Policy, page 10-3](#)
- [Message Splintering, page 10-5](#)
- [Configuring Mail Policies, page 10-7](#)

Overview of Mail Policies

The Email Security appliance enforces your organization's policies for messages sent to and from your users through the use of mail policies. These are sets of rules that specify the types of suspect, sensitive, or malicious content that your organization may not want entering or leaving your network. This content may include:

- spam
- legitimate marketing messages
- viruses
- phishing and other targeted mail attacks
- confidential corporate data
- personally identifiable information

You can create multiple policies that satisfy the disparate security needs of the different user groups within your organization. The Email Security appliance uses the rules defined in these policies to scan each message and, if necessary, perform an action to protect your user. For example, policies can prevent the delivery of suspected spam messages to executives while allowing their delivery to IT staff but with a modified subject to warn them of the content, or drop dangerous executable attachments for all users except those in the System Administrator group.

How to Enforce Mail Policies on a Per-User Basis

	Do This	More Info
Step 1	Enable the content-scanning features that you want the Email Security appliance to use for incoming or outgoing messages.	The features you can enable and configure one or more of the following: • Anti-Virus • File Reputation Filtering and File Analysis (incoming messages only) • Anti-Spam • Outbreak Filters • Data Loss Prevention (outgoing messages only) • Content Filters
Step 2	(Optional) Create content filters for actions to take on messages that contain specific data.	See Chapter 11, “Content Filters.”
Step 3	(Optional) Define an LDAP group query in order to specify users to whom the mail policy rules apply.	See Using Group LDAP Queries to Determine if a Recipient is a Group Member , page 25-23.
Step 4	(Optional) Define the default mail policies for incoming or outgoing messages.	See Configuring the Default Mail Policy for Incoming or Outgoing Messages , page 10-7.
Step 5	Define the group of users for whom you want to set up user-specific mail policies.	Create an incoming or outgoing mail policy. See Configuring Mail Policies, page 10-7 for more information.
Step 6	Configure the content security features and the content filter actions the appliance takes on messages.	Configure the different content security features for the mail policy. • Content Filters: Applying the Content Filter to Messages for a Certain User Group, page 11-18 • Anti-Virus: Configuring Virus Scanning Actions for Users, page 12-7 • File Reputation Filtering and File Analysis: Configuring the Incoming Mail Policy for File Reputation Scanning and File Analysis, page 16-6 • Anti-Spam: Defining Anti-Spam Policies, page 13-7 • Outbreak Filters: The Outbreak Filters Feature and Mail Policies, page 14-15 • Data Loss Prevention: Using Outgoing Mail Policies to Assign DLP Policies to Senders and Recipients, page 17-23 and About Associating Outgoing Mail Policies with DLP Policies in Enterprise Manager Deployments, page 17-31.

Handling Incoming and Outgoing Messages Differently

The Email Security appliances uses two different sets of mail policies for message content security:

- *Incoming mail policies* for messages are messages received from connections that match an ACCEPT HAT policy in any listener.
- *Outgoing mail policies* for messages are messages from connections that match a RELAY HAT policy in any listener. This includes any connection that was authenticated with SMTP AUTH.

Having separate sets of policies allow you to define different security rules for messages sent to your users and messages sent from your users. You manage these policies using the **Mail Policies > Incoming Mail Policies** or **Outgoing Mail Policies** pages in the GUI, or the `policyconfig` command in the CLI.

**Note**

Some features can be applied only to incoming or to outgoing mail policies. For example, Data Loss Prevention scanning can only be performed on outgoing messages. Advanced Malware Protection (File Reputation scanning and File Analysis) is available only in Incoming Mail Policies.

**Note**

In certain installations, “internal” mail being routed through the Cisco appliance may be considered *outgoing*, even if all the recipients are addressed to internal addresses. For example, by default for C170 appliances, the system setup wizard will configure only one physical Ethernet port with one listener for receiving inbound email and relaying outbound email.

Matching Users to a Mail Policy

As messages are received by the appliance, the Email Security appliance attempts to match each message recipient and sender to a mail policy in the Incoming or Outgoing Mail Policies table, depending on whether it is an incoming or outgoing message.

Matches are based on the recipient’s address, the sender’s address, or both:

- *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 mail policy.

- *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.

Related Topics

- [First Match Wins, page 10-4](#)
- [Examples of Policy Matching, page 10-4](#)

First Match Wins

Each user (sender or recipient) is evaluated for each mail policy defined the appropriate mail policy table in a top-down fashion.

For each user, the first matching policy wins. If a user does not match any specific policy, user will automatically match the default policy of the table.

If a match is made based on a sender address, all remaining recipients of a message will match that policy. (This is because there can be only one sender per message.)

Examples of Policy Matching

The following examples help show how the policy tables are matched in a top-down fashion.

Given the following Incoming Mail Email Security Policy table shown in [Table 10-1](#), incoming messages will match different policies.

Table 10-1 Policy Matching Example

Order	Policy Name	Users	
		Sender	Recipient
1	special_people	ANY	joe@example.com ann@example.com
2	from_lawyers	@lawfirm.com	ANY
3	acquired_domains	ANY	@newdomain.com @anotherexample.com
4	engineering	ANY	PublicLDAP ldapgroup: engineers
5	sales_team	ANY	jim@ john@ larry@
	Default Policy	ANY	ANY

Related Topics

- [Example 1, page 10-4](#)
- [Example 2, page 10-4](#)
- [Example 3, page 10-5](#)

Example 1

A message from sender `bill@lawfirm.com` sent to recipient `jim@example.com` will match policy #2, because the user description matches the sender (`@lawfirm.com`) and the recipient (ANY).

Example 2

Sender `joe@yahoo.com` sends an incoming message with three recipients: `john@example.com`, `jane@newdomain.com`, and `bill@example.com`:

- The message for recipient `jane@newdomain.com` will receive the anti-spam, anti-virus, outbreak filters, and content filters defined in policy #3.
- The message for recipient `john@example.com` will receive the settings defined in policy #5.
- Because the recipient `bill@example.com` does not match the engineering LDAP query, the message will receive the settings defined by the default policy.

This example shows how messages with multiple recipients can incur *message splintering*. See [Message Splintering, page 10-5](#) for more information.

Example 3

Sender `bill@lawfirm.com` sends a message to recipients `ann@example.com` and `larry@example.com`:

- The recipient `ann@example.com` will receive the anti-spam, anti-virus, outbreak filters, and content filters defined in policy #1.
- The recipient `larry@example.com` will receive the anti-spam, anti-virus, outbreak filters, and content filters defined in policy #2, because the sender (`@lawfirm.com`) and the recipient (`ANY`) matches.

Message Splintering

Intelligent message splintering is the mechanism that allows for differing recipient-based content security rules to be applied independently to message with multiple recipients.


Each recipient is evaluated for each policy in the appropriate mail policy 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, Advanced Malware Protection (incoming messages only), DLP scanning (outgoing messages only), Outbreak Filters, and content filters independently in the email pipeline.

[Table 10-2](#) illustrates the point at which messages are splintered in the email pipeline.

Table 10-2 **Message Splintering in the Email Pipeline**

**Note**

New MIDs (message IDs) are created for each message splinter (for example, MID 1 becomes MID 2 and MID 3). For more information, see the “Logging” chapter. In addition, the `trace` function shows which policies cause a message to be split.

Policy matching and message splintering in Email Security Manager policies obviously affect how you manage the message processing available on the appliance.

Related Topics

- [Managed Exceptions, page 10-6](#)

Managed Exceptions

Because the iterative processing of each splinter message impacts performance, Cisco recommends configuring your content security rules on a *managed exception* basis. In other words, evaluate your organization’s needs and try to configure the feature so that the majority of messages will be handled by the default mail policy and the minority of messages will be handled by a few additional “exception” policies. In this manner, message splintering will be minimized and you are less likely to impact system performance from the processing of each splinter message in the work queue.

Configuring Mail Policies

Mail policies map different user groups to specific security settings, such as Anti-Spam or Anti-Virus.

Related Topics

- [Configuring the Default Mail Policy for Incoming or Outgoing Messages, page 10-7](#)
- [Creating a Mail Policy for a Group of Senders and Recipients, page 10-7](#)
- [Finding Which Policies Apply to a Sender or Recipient, page 10-10](#)

Configuring the Default Mail Policy for Incoming or Outgoing Messages

The default mail policies apply to messages that are not covered by any other mail policy. If no other policies are configured, the default policies apply to all messages.

Before You Begin

Understand how you can define the individual security services for the mail policy. See [How to Enforce Mail Policies on a Per-User Basis, page 10-2](#).

Procedure

Step 1 Choose **Mail Policies > Incoming Mail Policies**

or

Mail Policies > Outgoing Mail Policies.

Step 2 Click the link for the security service you want to configure for the Default mail policy.



Note For default security service settings, the first setting on the page defines whether the service is enabled for the policy. You can click “Disable” to disable the service altogether.

Step 3 Configure the settings for the security service.

Step 4 Click **Submit**.

Step 5 Submit and commit your changes.

Creating a Mail Policy for a Group of Senders and Recipients

Before You Begin

- Understand how you can define the individual security services for the mail policy. See [How to Enforce Mail Policies on a Per-User Basis, page 10-2](#).
- Remember that each recipient is evaluated for each policy in the appropriate table (incoming or outgoing) in a top-down fashion. See [First Match Wins, page 10-4](#) for more information.

- (Optional) Define the delegated administrators who will be responsible for managing the mail policy. Delegated administrators can edit a policy's Anti-Spam, Anti-Virus, Advanced Malware Protection, and Outbreak Filters settings and enable or disable content filters for the policy. Only operators and administrators can modify a mail policy's name or its senders, recipients, or groups. Custom user roles that have full access to mail policies are automatically assigned to mail policies.

Procedure

-
- Step 1** Choose **Mail Policies > Incoming Mail Policies** or **Mail Policies > Outgoing Mail Policies**.
- Step 2** Click **Add Policy**.
- Step 3** Enter a name for the mail policy.
- Step 4** (Optional) Click the Editable by (Roles) link and select the custom user roles for the delegated administrators who will be responsible for managing the mail policy.
- Step 5** Define users for the policy. For instructions to define users, see [Defining Senders and Recipients for Mail Policies, page 10-8](#).
- Step 6** Click **Submit**.
- Step 7** Click the link for the content security service you want to configure for the mail policy.
- Step 8** From the drop-down list, select the option to customize the settings for the policy instead of using the default settings.
- Step 9** Customize the security service settings.
- Step 10** Submit and commit your changes.
-

Related Topics

- [How to Configure the Appliance to Scan Messages for Spam, page 13-2](#)
- [Defining Senders and Recipients for Mail Policies, page 10-8](#)

Defining Senders and Recipients for Mail Policies

You can define senders and recipients to whom the policy applies in the following ways:

- Full email address: `user@example.com`
- Partial email address: `user@`
- All users in a domain: `@example.com`
- All users in a partial domain: `@.example.com`
- By matching an LDAP Query



Note

Entries for users are case-insensitive in both the GUI and CLI in AsyncOS. For example, if you enter the recipient `Joe@` for a user, a message sent to `joe@example.com` will match.

While defining senders and recipients for mail policies, keep in mind that:

- You must specify at least one sender and recipient.
- You can set the policy to match if,
 - The message is from any sender, one or more of the specified senders, or none of the specified senders.
 - The message is sent to any recipient, one or more of the specified recipients, or all of the specified recipients and none of the specified recipients.

Procedure

Step 1 Under **Users** section, click **Add User**.

Step 2 Define the senders for the policy. Choose one of the following options:

- **Any Sender.** The policy is matched if the message is from any sender.
- **Following Senders.** The policy is matched if the message is from one or more of the specified senders. Select this option and enter sender details in the text box or choose an LDAP group query.
- **Following Senders are Not.** The policy is matched if the message is not from none of the specified senders. Select this option and enter sender details in the text box or choose an LDAP group query.

To understand how sender conditions are set while choosing the above fields, see [Examples, page 10-10](#).

Step 3 Define the recipients for the policy. Choose one of the following options:

- **Any Recipient.** The policy is matched if the message is sent to any recipient.
- **Following Recipients.** The policy is matched if the message is sent to the specified recipients. Select this option, enter the recipient details in the text box or choose an LDAP group query.
You can choose whether policy is matched if the message is sent to one or more of the specified recipients or all of the specified recipients. Choose one of the following options from the drop-down list: **If one more conditions match** or **Only if all conditions match**.
- **Following Recipients are Not.** The policy is matched if the message is sent to none of the specified recipients. Select this option, enter the recipient details in the text box or choose an LDAP group query.



Note You can configure this option only if you have selected **Following Recipients** and chosen **Only if all conditions match** from the drop-down list.

To understand how recipient conditions are set while choosing the above fields, see [Examples, page 10-10](#).

Step 4 Click **Submit**.

Step 5 Review the selected conditions on the **Users** section.

Related Topics

- [Creating a Mail Policy for a Group of Senders and Recipients, page 10-7](#)
- [Examples, page 10-10](#)

Examples

The following table describes how conditions are set when you choose various options on the Add User page.

Sender			Recipient			Condition
Any Sender	Following Senders	Following Senders are Not	Any Recipient	Following Recipients	Following Recipients are Not	
Selected	-	-	-	Selected (Default) Only if all conditions match option is selected Values: user1@, user2@	-	Sender: Any Recipient: user1@[AND]user2@
-	Selected Values: u1@a.com, u2@a.com	-	-	Selected (Default) Only if all conditions match option is selected Values: u1@b.com, u2@b.com	Selected Values: u3@b.com, u4@b.com	Sender: u1@a.com[OR]u2@a.com Recipient: [u1@b.com[AND]u2@b.com] [AND] [[NOT] [u3@b.com[AND]u4@b.com]]
-	-	Selected Values: u1@a.com, u2@a.com	-	Selected If one or more conditions match option is also selected Values: u1@b.com, u2@b.com	-	Sender: [NOT] [u1@a.com[OR]u2@a.com] Recipient: u1@b.com[OR]u2@b.com

Related Topics

- [Defining Senders and Recipients for Mail Policies, page 10-8](#)

Finding Which Policies Apply to a Sender or Recipient

Use the Find Policies section at the top of the Mail Policies page to search for users already defined in incoming or outgoing mail policies.

For example, type `bob@example.com` and click the Find Policies button to display results showing which policies contain defined users that will match the policy.

Click the name of the policy to edit the users for that policy.

Note that the default policy will always be shown when you search for any user, because, by definition, if a sender or recipient does not match any other configured policies, it will *always* match the default policy.

Related Topics

- [Managed Exceptions, page 10-11](#)

Managed Exceptions

Using the steps shown in the two examples above, you can begin to create and configure policies on a *managed exception* basis. In other words, after evaluating your organization's needs you can configure policies so that the majority of messages will be handled by the default policy. You can then create additional "exception" policies for specific users or user groups, managing the differing policies as needed. In this manner, message splintering will be minimized and you are less likely to impact system performance from the processing of each splinter message in the work queue.

You can define policies based on your organizations' or users' tolerance for spam, viruses, and policy enforcement. [Table 10-3 on page 10-11](#) outlines several example policies. "Aggressive" policies are designed to minimize the amount of spam and viruses that reach end-users mailboxes. "Conservative" policies are tailored to avoid false positives and prevent users from missing messages, regardless of policies.

Table 10-3 **Aggressive and Conservative Email Security Manager Settings**

	Aggressive Settings	Conservative Settings
Anti-Spam	Positively identified spam: Drop Suspected spam: Quarantine Marketing mail: Deliver and prepend "[Marketing]" to the subject messages	Positively identified spam: Quarantine Suspected spam: Deliver and prepend "[Suspected Spam]" to the subject of messages Marketing mail: Disabled
Anti-Virus	Repaired messages: Deliver Encrypted messages: Drop Unscannable messages: Drop Infectious messages: Drop	Repaired messages: Deliver Encrypted messages: Quarantine Unscannable messages: Quarantine Infectious messages: Drop
Advanced Malware Protection (File Reputation Filtering and File Analysis)	Unscanned attachments: Drop Messages with Malware Attachments: Drop Messages with pending File Analysis: Quarantine	Unscanned attachments: Deliver and prepend "[WARNING: ATTACHMENT UNSCANNED]" to the subject of messages. Messages with Malware Attachments: Drop Messages with pending File Analysis: Deliver and prepend "[WARNING: ATTACHMENT(S) MAY CONTAIN MALWARE]" to the subject of messages.
Virus Filters	Enabled, no specific filename extensions or domains allowed to bypass Enable message modification for all messages	Enabled with specific filename extensions or domains allowed to bypass Enable message modification for unsigned messages



Content Filters

- [Overview of Content Filters, page 11-1](#)
- [How Content Filters Work, page 11-1](#)
- [Filtering Messages Based on Content, page 11-16](#)

Overview of Content Filters

Sometimes the Email Security appliance receives a message that should be given special treatment due to its content, whether it's because the content warrants quarantining for later examination, because corporate policy requires certain messages to be encrypted before delivery, or any number of reasons. These are cases when the message cannot be handled by the other content security features, such as anti-virus scanning or DLP. The appliance uses content filters to scan for such content and then take appropriate action on the message.

How Content Filters Work

Content filters are similar to message filters, but are applied after the message has undergone message filters processing and anti-spam and anti-virus scanning. The Email Security appliance uses content filters to scan messages on a per-user (sender or recipient) basis. Content filters are similar to message filters, except that they are applied later in the email pipeline — after a message has been “splintered” into a number of separate messages for each matching mail policy. (See [Message Splintering, page 10-5](#) for more information.) The functionality of content filters is applied after message filters processing and anti-spam and anti-virus scanning have been performed on a message.

A content filter is limited to scanning either incoming or outgoing messages. You cannot define a filter that scans both types of messages. The Email Security appliance has a separate “master list” of content filters for each type of message. The master list also determines in which order the appliance runs the content filters. However, each individual mail policy determines which particular filters will be executed when a message matches the policy.

AsyncOS provides a “rule builder” page that allows you to easily create the content filter using four parts:

- *conditions* that trigger when the appliance uses a content filter to scan a message (optional)
- *actions* that the appliance takes on a message (required)
- *action variables* that the appliance can add to a message when modifying it (optional)

Related Topics

- [How to Scan Message Content Using a Content Filter, page 11-2](#)
- [Content Filter Conditions, page 11-2](#)
- [Content Filter Actions, page 11-9](#)
- [Action Variables, page 11-14](#)

How to Scan Message Content Using a Content Filter

	Do This	More Info
Step 1	(Optional) Define the supporting features for the content filter.	Create any of the following items that you want to use with your content filter: • Encryption Profile • Disclaimer template • Notification template • Policy quarantine • URL whitelists
Step 2	Define the incoming or outgoing content filter.	A content filter may comprise of three parts: • Content Filter Conditions (optional) • Content Filter Actions • Action Variables (optional) Creating a Content Filter, page 11-16
Step 3	Define the group of users for whom you want to set up content security rules.	Create an incoming or outgoing mail policy.
Step 4	Assign the content filter to the group of user whose incoming or outgoing messages you want to use the filter for.	See Chapter 10, “Mail Policies.”

Content Filter Conditions

A condition is a “trigger” that determines whether the Email Security appliance uses the filter on a message that matches the associated mail policy. Specifying conditions for a content filter is optional. Content filters without a condition are applied to all messages that match the associated mail policy.

In the content filter conditions, when you add filter rules that search for certain patterns in the message body or attachments, you can specify the minimum threshold for the number of times the pattern must be found. When AsyncOS scans the message, it totals the “score” for the number of matches it finds in the message and attachments. If the minimum threshold is not met, the regular expression does not evaluate to true. You can specify this threshold for text, smart identifiers, or content dictionary terms.

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”).

Table 11-1 **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	Contains text: Does the message body contain text or an attachment that matches a specific pattern? Contains smart identifier: Does content in the message body or attachment match a smart identifier? Contains term in content dictionary: Does the message body contain any of the regular expressions or terms in the content dictionary named <i><dictionary name></i>? For this option to be enabled, the dictionary must already have been created. See Content Dictionaries, page 21-2. Note The dictionary-related conditions are only available if you have one or more dictionaries enabled. For information about creating content dictionaries, see Content Dictionaries, page 21-2. Number of matches required. Specify the number of matches required for the rule to evaluate to true. You can specify this threshold for text, smart identifiers, or content dictionary terms. This includes delivery-status parts and associated attachments.

Table 11-1 **Content Filter Conditions (continued)**

Condition	Description
Message Body	Contains text: Does the message body contain text that matches a specific pattern? Contains smart identifier: Does content in the message body match a smart identifier? 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 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>? For this option to be enabled, the dictionary must already have been created. See Content Dictionaries, page 21-2. Note The dictionary-related conditions are only available if you have one or more dictionaries enabled. For information about creating content dictionaries, see Content Dictionaries, page 21-2. Number of matches required. Specify the number of matches required for the rule to evaluate to true. You can specify this threshold for text or smart identifiers. This rule applies to the body of the message only. It does not include attachments or headers.
URL Reputation	See Taking Action Based on the Reputation or Category of URLs in Messages, page 15-7 and Creating Whitelists for URL Filtering, page 15-4. Use "No Score" for URLs for which a reputation cannot be determined.
URL Category	See Filtering by URL Reputation or URL Category: Conditions and Rules, page 15-8 and About URL Categories, page 15-13 .
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.

Table 11-1 **Content Filter Conditions (continued)**

Condition	Description
Attachment Content	Contains text. Does the message contain an attachment that contains text or another attachment that matches a specific pattern? This rule is similar to the <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. Contains a smart identifier. Does content in the message attachment match the specified smart identifier? Contains terms in content dictionary. Does the attachment contain any of the regular expressions or terms in the content dictionary named <i><dictionary name></i>? To search for dictionary terms, the dictionary must already have been created. See Content Dictionaries, page 21-2. Note The dictionary-related conditions are only available if you have one or more dictionaries enabled. For information about creating content dictionaries, see Content Dictionaries, page 21-2. Number of matches required. Specify the number of matches required for the rule to evaluate to true. You can specify this threshold for text, smart identifier, or content dictionary matches.

Table 11-1 Content Filter Conditions (continued)

Condition	Description
Attachment File Info	Filename. Does the message have an attachment with a filename that matches a specific pattern?
	Filename contains term in content dictionary. Does the message have an attachment with a filename that contains any of the regular expressions or terms in the content dictionary named <i><dictionary name></i> ?
	For this option to be enabled, the dictionary must already have been created. See Content Dictionaries, page 21-2 .
	Note The dictionary-related conditions are only available if you have one or more dictionaries enabled. For information about creating content dictionaries, see Content Dictionaries, page 21-2 .
	File type. Does the message have an attachment of a file type that matches a specific pattern based on its fingerprint (similar to a UNIX <i>file</i> command)?
	MIME type. Does the message have an attachment of a specific MIME type? This rule is similar to the <i>attachment-type</i> 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.)
Attachment Protection	Image Analysis. Does the message have an image attachment that matches the image verdict specified? Valid image analysis verdicts include: <i>Suspect, Inappropriate, Suspect or Inappropriate, Unscannable, or Clean.</i>
	Attachment is Corrupt. Does this message have an attachment that is corrupt?
	Note A corrupt attachment is an attachment that the scanning engine cannot scan and identified as corrupt.
Attachment Protection	Contains an attachment that is password-protected or encrypted. (For example, use this condition to identify attachments that are potentially unscannable.)
	Contains an attachment that is NOT password-protected or encrypted.

Table 11-1 Content Filter Conditions (continued)

Condition	Description
Subject Header	Subject Header: Does the subject header match a certain pattern? Contains terms in content dictionary: Does the subject header contain any of the regular expressions or terms in the content dictionary <i><dictionary name></i>? To search for dictionary terms, the dictionary must already have been created. See Content Dictionaries, page 21-2. Note The dictionary-related conditions are only available if you have one or more dictionaries enabled. For information about creating content dictionaries, see Content Dictionaries, page 21-2.
Other Header	Header name: Does the message contain a specific header? Header value: Does the value of that header match a certain pattern? Header value contains terms in the content dictionary. Does the specified header contain any of the regular expressions or terms in the content dictionary named <i><dictionary name></i>? To search for dictionary terms, the dictionary must already have been created. See Content Dictionaries, page 21-2. Note The dictionary-related conditions are only available if you have one or more dictionaries enabled. For information about creating content dictionaries, see Content Dictionaries, page 21-2. For an example showing how this option can be used, see Using Custom Headers to Redirect URLs in Suspected Spam to the Cisco Web Security Proxy: Configuration Example, page 13-11.

Table 11-1 Content Filter Conditions (continued)

Condition	Description
Envelope Sender	Envelope Sender. Does the Envelope Sender (i.e., the Envelope From, <MAIL FROM>) match a given pattern? Matches LDAP group. Is the Envelope Sender, i.e., the Envelope From, <MAIL FROM>) in a given LDAP group? Contains term in content dictionary. Does the envelope sender contain any of the regular expressions or terms in the content dictionary named <dictionary name>? To search for dictionary terms, the dictionary must already have been created. See Content Dictionaries, page 21-2. Note The dictionary-related conditions are only available if you have one or more dictionaries enabled. For information about creating content dictionaries, see Content Dictionaries, page 21-2.
Envelope Recipient	Envelope Recipient. Does the Envelope Recipient, (i.e. the Envelope To, <RCPT TO>) match a given pattern? Matches LDAP group. Is the Envelope Recipient, (i.e. the Envelope To, <RCPT TO>) in a given LDAP group? Contains term in content dictionary. Does the envelope recipient contain any of the regular expressions or terms in the content dictionary named <dictionary name>? To search for dictionary terms, the dictionary must already have been created. See Content Dictionaries, page 21-2. Note The dictionary-related conditions are only available if you have one or more dictionaries enabled. For information about creating content dictionaries, see Content Dictionaries, page 21-2. Note: The Envelope Recipient rule is message-based. If a message has multiple recipients, only one recipient has to be found in a group for the specified action to affect the message to all recipients. Is the Envelope Sender (i.e., the Envelope From, <MAIL FROM>) in a given LDAP group?
Receiving Listener	Did the message arrive via the named listener? The listener name must be the name of a listener currently configured on the system.

Table 11-1 **Content Filter Conditions (continued)**

Condition	Description
Remote IP	Was the message sent from a remote host that matches a given IP address or IP block? The Remote IP rule tests to see if the IP address of the host that sent that message matches a certain pattern. This can be an Internet Protocol version 4 (IPv4) or version 6 (IPv6) address. The IP address pattern is specified using the allowed hosts notation described in Sender Group Syntax, page 7-4 , except for the SBO, SBRS, dnslist notations and the special keyword ALL.
Reputation Score	What is the sender's SenderBase Reputation Score? The Reputation Score rule checks the SenderBase Reputation Score against another value.
DKIM Authentication	Did DKIM authentication pass, partially verify, return temporarily unverifiable, permanently fail, or were no DKIM results returned?
SPF Verification	What was the SPF verification status? This filter rule allows you to query for different SPF verification results. For more information about SPF verification, see the "Email Authentication" chapter.
S/MIME Gateway Message	Is the message S/MIME signed, encrypted, or signed and encrypted? For more information, see Chapter 19, "S/MIME Security Services."
S/MIME Gateway Verified	Is the S/MIME message successfully verified, decrypted, or decrypted and verified? For more information, see Chapter 19, "S/MIME Security Services."

Content Filter Actions

The action is what the Email Security appliance does with a message that matches the content filter's condition. Many different types of actions are available, including modifying the message, quarantining it, or dropping it. A "final action" performed on a message, delivering or dropping it, forces the Email Security appliance to perform the action immediately and forgo all further processing, such as Outbreak Filter or DLP scanning.

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 Text Resources chapter.

Only one final action may be defined per filter, and the final action must be last action listed. Bounce, deliver, and drop are final actions. When entering actions for content filters, the GUI and CLI will force final actions to be placed last.

Table 11-2 **Content Filter Actions**

Action	Description
Quarantine	Quarantine. Flags the message to be held in one of the policy quarantine areas. Duplicate message: Sends a copy of the message to the specified quarantine and continues processing the original message. Any additional actions apply to the original message.
Encrypt on Delivery	The message continues to the next stage of processing. When all processing is complete, the message is encrypted and delivered. Encryption rule: Always encrypts the message or only encrypts it if an attempt to send it over a TLS connection first fails. See Using a TLS Connection as an Alternative to Encryption, page 18-9 for more information. Encryption Profile. Once processing is complete, encrypts the message using the specified encryption profile, then delivers the message. This action is for use with a Cisco Encryption Appliance or a hosted key service. Subject. Subject for the encrypted message. By default, the value is <code>\$Subject</code>.
Strip Attachment by Content	Attachment contains. Drops all attachments on messages that contain the regular expression. Archive files (zip, tar) will be dropped if any of the files they contain match the regular expression pattern. Contains smart identifier. Drops all attachments on a message that contains the specified smart identifier. Attachment contains terms in the content dictionary. Does the attachment contain any of the regular expressions or terms in the content dictionary named <code><dictionary name>?</code> Number of matches required. Specify the number of matches required for the rule to evaluate to true. You can specify this threshold for text, smart identifier, or content dictionary matches. Replacement message. The optional comment serves as the means to modify the text used to replace the attachment that was dropped. Attachment footers simply append to the message.

Table 11-2 Content Filter Actions (continued)

Action	Description
Strip Attachment by File Info	File name. Drops all attachments on messages that have a filename that match the given regular expression. Archive file attachments (zip, tar) will be dropped if they contain a file that matches. File size. Drops all attachments on the message that, in raw encoded form, are equal to or greater than the size (in bytes) given. Note that for archive or compressed files, this action does not examine the uncompressed size, but rather the size of the actual attachment itself. File type. Drops all attachments on messages that match the given “fingerprint” of the file. Archive file attachments (zip, tar) will be dropped if they contain a file that matches. MIME type. Drops all attachments on messages that have a given MIME type. Image Analysis Verdict. Drops attachments for image attachments that match the image verdict specified. Valid image analysis verdicts include: <i>Suspect</i>, <i>Inappropriate</i>, <i>Suspect or Inappropriate</i>, <i>Unscannable</i>, or <i>Clean</i>. 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.
URL Reputation	See Modifying URLs in Messages: Using URL Reputation and URL Category Actions in Filters, page 15-8 and Creating Whitelists for URL Filtering, page 15-4. Use "No Score" to specify an action for URLs for which a reputation cannot be determined.
URL Category	See Modifying URLs in Messages: Using URL Reputation and URL Category Actions in Filters , page 15-8 and About URL Categories , page 15-13.
Add Disclaimer Text	Above. Add disclaimer above message (heading). Below. Add disclaimer below message (footer). Note: You must have already created disclaimer text in order to use this content filter action. See Disclaimer Template, page 21-12 for more information.
Bypass Outbreak Filter Scanning	Bypass Outbreak Filter scanning for this message.

Table 11-2 Content Filter Actions (continued)

Action	Description
Bypass DKIM Signing	Bypass DKIM signing for this message.
Send Copy (Bcc:)	Email addresses. Copies the message anonymously to the specified recipients. Subject. Add a subject for the copied message. Return path (optional). Specify a return path. Alternate mail host (optional). Specify an alternate mail host.
Notify	Notify. Reports this message to the specified recipients. You can optionally notify the sender and recipients. Subject. Add a subject for the copied message. Return path (optional). Specify a return path. Use template. Select a template from the templates you created. Include original message as an attachment. Adds the original message as an attachment.
Change Recipient to	Email address. Changes the recipient of the message to the specified email address.
Send to Alternate Destination Host	Mail host. Changes the destination mail host for the message to the specified mail host. Note This action prevents a message classified as spam by an anti-spam scanning engine from being quarantined. This action overrides the quarantine and sends it to the specified mail host.
Deliver from IP Interface	Send from IP interface. Send from the specified IP Interface. The Deliver from IP Interface action changes the source host for the message to the source specified. The source host consists of the IP interface that the messages should be delivered from.
Strip Header	Header name. Remove the specified header from the message before delivering.

Table 11-2 Content Filter Actions (continued)

Action	Description
Add/Edit Header	Inserts a new header into the message or modifies an existing header. Header name. Name of new or existing header. Specify value of new header. Inserts a value for the new header into the message before delivering. Prepend to the Value of Existing Header. Prepends the value to the existing header before delivering. Append to the Value of Existing Header. Appends the value to the existing header before delivering. Search & Replace from the Value of Existing Header. Enter a search term to find the value you want to replace in the existing header in the Search for field. Enter the value you want to insert into the header in the Replace with field. You can use a regular expression to search for the value. Leave the Replace with field empty if you want to delete the value from the header.
Add Message Tag	Inserts a custom term into the message to use with RSA Email DLP policy filtering. You can configure a RSA Email DLP policy to limit scanning to messages with the message tag. The message tag is not visible to recipients. For information on using messages tags in a DLP policy, see DLP Policies for RSA Email DLP, page 17-6.
Add Log Entry	Inserts customized text into the IronPort Text Mail logs at the <code>INFO</code> level. The text can include action variables. The log entry also appears in message tracking.
S/MIME Sign/Encrypt on Delivery	Performs an S/MIME signing or encryption of the message during the delivery. This means that the message continues to the next stage of processing, and when all processing is complete, the message is signed or encrypted and delivered. S/MIME Sending Profile: Performs an S/MIME signing or encryption using the specified S/MIME sending profile. See Managing S/MIME Sending Profiles, page 19-10.

Table 11-2 Content Filter Actions (continued)

Action	Description
Encrypt and Deliver Now (Final Action)	Encrypts and delivers the message, skipping any further processing. Encryption rule: Always encrypts the message or only encrypts it if an attempt to send it over a TLS connection first fails. See Using a TLS Connection as an Alternative to Encryption, page 18-9 for more information. Encryption Profile. Encrypts the message using the specified encryption profile, then delivers the message. This action is for use with a Cisco Encryption Appliance or a hosted key service. Subject. Subject for the encrypted message. By default, the value is \$Subject.
S/MIME Sign/Encrypt (Final Action)	Performs an S/MIME signing or encryption and delivers the message, skipping any further processing. S/MIME Sending Profile: Performs an S/MIME signing or encryption using the specified S/MIME sending profile. See Managing S/MIME Sending Profiles, page 19-10.
Bounce (Final Action)	Sends the message back to the sender.
Skip Remaining Content Filters (Final Action)	Delivers the message to the next stage of processing, skipping any further content filters. Depending on configuration, this may mean deliver the message to recipient(s), quarantine, or begin Outbreak Filters scanning.
Drop (Final Action)	Drops and discards the message.

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 appliance supports the following set of action variables:

Table 11-3 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.

Table 11-3 Action Variables (continued)

Variable	Syntax	Description
Envelope Sender	<code>\$envelopefrom</code> or <code>\$envelopesender</code>	Replaced by the Envelope Sender (Envelope From, <MAIL FROM>) of the message.
Envelope Recipients	<code>\$EnvelopeRecipients</code>	Replaced by all Envelope Recipients (Envelope To, <RCPT TO>) of the message.
File Names	<code>\$filenames</code>	Replaced with a comma-separated list of the message's attachments' filenames.
File Sizes	<code>\$filesizes</code>	Replaced with a comma-separated list of the message's attachment's file sizes.
File Types	<code>\$filetypes</code>	Replaced with a comma-separated list of the message's attachments' file types.
Filter Name	<code>\$FilterName</code>	Replaced by the name of the filter being processed.
GMTimeStamp	<code>\$GMTimeStamp</code>	Replaced by the current time and date, as would be found in the Received: line of an email message, using GMT.
HAT Group Name	<code>\$Group</code>	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	<code>\$Policy</code>	Replaced by the name of the HAT policy applied to the sender when injecting the message. If no predefined policy name was used, the string ">Unknown<" is inserted.
Matched Content	<code>\$MatchedContent</code>	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	<code>\$Header['string']</code>	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	<code>\$Hostname</code>	Replaced by the hostname of the Email Security appliance.
Internal Message ID	<code>\$MID</code>	Replaced by the Message ID, or "MID" used internally to identify the message. Not to be confused with the RFC822 "Message-Id" value (use <code>\$Header</code> to retrieve that).
Receiving Listener	<code>\$RecvListener</code>	Replaced by the nickname of the listener that received the message.
Receiving Interface	<code>\$RecvInt</code>	Replaced by the nickname of the interface that received the message.

Table 11-3 **Action Variables (continued)**

Variable	Syntax	Description
Remote IP Address	<code>\$RemoteIP</code>	Replaced by the IP address of the system that sent the message to the Email Security appliance.
Remote Host Address	<code>\$remotehost</code>	Replaced by the hostname of the system that sent the message to the appliance.
SenderBase Reputation Score	<code>\$Reputation</code>	Replaced by the SenderBase Reputation score of the sender. If there is no reputation score, it is replaced with “None”.
Subject	<code>\$Subject</code>	Replaced by the subject of the message.
Time	<code>\$Time</code>	Replaced by the current time, in the local time zone.
Timestamp	<code>\$Timestamp</code>	Replaced by the current time and date, as would be found in the Received: line of an email message, in the local time zone.

Filtering Messages Based on Content

Related Topics

- [Creating a Content Filter, page 11-16](#)
- [Enabling Content Filters for All Recipients by Default, page 11-17](#)
- [Applying the Content Filter to Messages for a Certain User Group, page 11-18](#)
- [Notes on Configuring Content Filters in the GUI, page 11-18](#)

Creating a Content Filter

Before You Begin

- If you want to encrypt a message that matches the content filter, create an encryption profile.
- If you want to add a disclaimer to a matching message, create a disclaimer template to use for generating disclaimers.
- If you want to send a notification message to a user due to a matching message, create a notification template for generating notifications.
- If you want to quarantine a message, you create a new policy quarantine for these messages or use an existing one.

Procedure

Step 1 Click **Mail Policies > Incoming Mail Policies**

or

Mail Policies > Outgoing Mail Policies.

- Step 2** Click **Add Filter**.
- Step 3** Enter a name and description for the filter.
- Step 4** (X-REF) Click the **Editable By (Roles)** link, select the Policy Administrator and click **OK**.
Delegated administrators who belong to the Policy Administrator user role will be able to edit this content filter and use it in their mail policies.
- Step 5** (Optional) Add a condition for triggering the filter.
- Click **Add Condition**.
 - Select the condition type.
 - Define the condition's rules.
 - Click **OK**.
 - Repeat these steps for any additional conditions you want to add to the filter. When you define more than one condition for a content filter, you can define whether *all* of the defined actions (that is, a logical AND) or *any* of the defined actions (logical OR) need to apply in order for the content filter to be considered a match.



Note If you do not add a condition, the appliance will perform the content filter's action to any message that matches one of the mail policies associated with the filter.

- Step 6** Add an action for the appliance to take on a message that matches the filter's condition.
- Click **Add Action**.
 - Select the action type.
 - Define the action.
 - Click **OK**.
 - Repeat the previous steps for any additional actions you want the appliance to take.
 - For multiple actions, arrange the actions in the order that you want the appliance to apply them to the message. There can only be one "final" action per filter, and AsyncOS automatically moves the final action to the end of the order.
- Step 7** Submit and commit your changes.

What to Do Next

- You can enable the content filter in a default incoming or outgoing mail policy.
- You can enable the content filter in a mail policy for a specific group of users.

Enabling Content Filters for All Recipients by Default

Procedure

- Step 1** Click **Mail Policies > Incoming Mail Policies**
or
Mail Policies > Outgoing Mail Policies.

- Step 2** Click the link for the Content Filters security service in the default policy row.
- Step 3** On the Content Filtering security service page, change the value Content Filtering for Default Policy from “Disable Content Filters” to “Enable Content Filters (Customize settings).”
- The content filters defined in the master list (which were created in [Overview of Content Filters, page 11-1](#)) are displayed on this page. When you change the value to “Enable Content Filters (Customize settings),” the checkboxes for each filter become enabled.
- Step 4** Check the **Enable** checkbox for each content filter you want to enable.
- Step 5** Submit and commit your changes.
-

Applying the Content Filter to Messages for a Certain User Group

Before You Begin

- Create an incoming or outgoing mail policy for the group of users whose messages for which you want to use the content filter. See [Creating a Mail Policy for a Group of Senders and Recipients, page 10-7](#) for more information.

Procedure

-
- Step 1** Click **Mail Policies > Incoming Mail Policies**
- or
- Mail Policies > Outgoing Mail Policies.**
- Step 2** Click the link for the Content Filters security service (the Content Filters column) for the mail policy to which you want to apply the content filter.
- Step 3** On the Content Filtering security service page, change the value for Content Filtering for Policy: Engineering from “Enable Content Filtering (Inherit default policy settings)” to “Enable Content Filtering (Customize settings).”
- Step 4** Select the checkboxes for the content filters you want to use.
- Step 5** Submit and commit your changes.
-

Notes on Configuring Content Filters in the GUI

- It is not necessary to specify a condition when creating a content filter. When no action is defined, any actions defined will always apply in the rule. (Specifying no condition is equivalent to using the `true()` message filter rule — all messages will be matched if the content filter is applied to a policy.)
- If you do not assign a custom user role to a content filter, the content filter is public and can be used by any delegated administrator for their mail policies. See the “Common Administrative Tasks” chapter for more information on delegated administrators and content filters.
- Administrators and operators can view and edit all content filters on an appliance, even when the content filters are assigned to custom user roles.
- When entering text for filter rules and actions, the following meta characters have special meaning in regular expression matching: `. ^ $ * + ? { [ ] \ | ( )`

If you do not wish to use regular expression you should use a '\' (backslash) to escape any of these characters. For example: "*Warning*"

- You can test message splintering and content filters by creating “benign” content filters. For example, it is possible to create a content filter whose only action is “deliver.” This content filter will not affect mail processing; however, you can use this filter to test how Email Security Manager policy processing affects other elements in the system (for example, the mail logs).
- Conversely, using the “master list” concept of the Incoming or Outgoing Content Filters, it is possible to create very powerful, wide-sweeping content filters that will immediately affect message processing for all mail handled by the appliance. The process for this is to:
 - Use the Incoming or Outgoing Content Filters page to create a new content filter whose order is 1.
 - Use the Incoming or Outgoing Mail Policies page to enable the new content filter for the default policy.
 - Enable the content filter for all remaining policies.
- The Bcc: and Quarantine actions available in Content Filters can help you determine the retention settings of quarantines you create. (See [Chapter 30, “Policy, Virus, and Outbreak Quarantines.”](#)) You can create filters that would simulate mail flow into and out of your policy quarantines so that messages are not released too quickly from the system (that is, the quarantine areas do not fill their allotted disk space too quickly).
- Because it uses the same settings as the Scan Behavior page or the `scanconfig` command, the “Entire Message” condition does not scan a message’s headers; choosing the “Entire Message” will scan only the message body and attachments. Use the “Subject” or “Header” conditions to search for specific header information.
- Configuring users by LDAP query will only appear in the GUI if you have LDAP servers configured on the appliance (that is, you have configured the appliance to query specific LDAP servers with specific strings using the `ldapconfig` command).
- Some sections of the content filter rule builder will not appear in the GUI if the resource has not been preconfigured. For example, notification templates and message disclaimers will not appear as options if they have not been configured previously using the Text Resources page or the `textconfig` command in the CLI.
- Content filters features will recognize, can contain, and/or scan for text in the following character encodings:
 - Unicode (UTF-8)
 - Unicode (UTF-16)
 - Western European/Latin-1 (ISO 8859-1)
 - Western European/Latin-1 (Windows CP1252)
 - Traditional Chinese (Big 5)
 - Simplified Chinese (GB 2312)
 - Simplified Chinese (HZ GB 2312)
 - Korean (ISO 2022-KR)
 - Korean (KS-C-5601/EUC-KR)
 - Japanese (Shift-JIS (X0123))
 - Japanese (ISO-2022-JP)
 - Japanese (EUC)

You can mix and match multiple character sets within a single content filter. Refer to your web browser's documentation for help displaying and entering text in multiple character encodings. Most browsers can render multiple character sets simultaneously.

- On the Incoming or Outgoing Content Filters summary pages, use the links for “Description,” “Rules,” and “Policies” to change the view presented for the content filters:
 - The **Description** view shows the text you entered in the description field for each content filter. (This is the default view.)
 - The **Rules** view shows the rules and regular expressions build by the rule builder page.
 - The **Policies** shows the policies for which each content filter is enabled.

Figure 11-1 Using the Links to Toggle Description, Rules, and Policy for Content Filters Incoming Content Filters

Filters				
Add Filter...				
Order	Filter Name	Description Rules Policies	Duplicate	Delete
1	scan_for_confidential	scan_for_confidential: if (body-contains("confidential")) { quarantine ("Policy"); bcc ("hr@example.com", "[message matched confidential filter]"); }		
2	no_mp3s	no_mp3s: if (true) { drop-attachments-by-filetype("mp3", "mp3 deleted"); }		
3	ex_employee	ex_employee: if (rcpt-to == "^doug@") { notify-copy ("{\$EnvelopeSender}", "message bounced for ex-employee of example.com"); bounce(); }		
4	drop_large_attachments	drop_large_attachments: if (true) { drop-attachments-by-size(5242880, "This attachment was too big!"); }		



Anti-Virus

- [Anti-Virus Scanning Overview, page 12-1](#)
- [Sophos Anti-Virus Filtering, page 12-2](#)
- [McAfee Anti-Virus Filtering, page 12-5](#)
- [How to Configure the Appliance to Scan for Viruses, page 12-6](#)
- [Sending an Email to the Appliance to Test Anti-Virus Scanning, page 12-16](#)
- [Updating Virus Definitions, page 12-18](#)

Anti-Virus Scanning Overview

The Cisco appliance includes integrated virus scanning engines from third party companies Sophos and McAfee. You can obtain license keys for the Cisco appliance to scan messages for viruses using one or both of these virus scanning engines, and then configure your appliance to scan for viruses using either anti-virus scanning engine.

The McAfee and Sophos engines contain the program logic necessary to scan files at particular points, process and pattern-match virus definitions with data they find in your files, decrypt and run virus code in an emulated environment, apply heuristic techniques to recognize new viruses, and remove infectious code from legitimate files.

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

If enabled, virus scanning is performed in the “work queue” on the appliance, immediately after Anti-Spam scanning. (See [Email Pipeline and Security Services, page 4-7.](#))

By default, virus scanning is enabled for the default incoming and outgoing mail policies.

Related Topics

- [Evaluation Key, page 12-2](#)
- [Scanning Messages with Multiple Anti-Virus Scanning Engines, page 12-2](#)

Evaluation Key

Your Cisco appliance ships with a 30-day evaluation key for each available anti-virus scanning engine. You enable the evaluation key by accessing the license agreement in the System Setup Wizard or Security Services > Sophos/McAfee Anti-Virus pages (in the GUI) or running the `antivirusconfig` or `systemsetup` commands (in the CLI). Once you have accepted the agreement, the Anti-Virus scanning engine will be enabled, by default, for the default incoming and outgoing mail policies. For information on enabling the feature beyond the 30-day evaluation period, contact your Cisco sales representative. You can see how much time remains on the evaluation via the System Administration > Feature Keys page or by issuing the `featurekey` command. (For more information, see [Feature Keys, page 33-5](#).)

Scanning Messages with Multiple Anti-Virus Scanning Engines

AsyncOS supports scanning messages with multiple anti-virus scanning engines — multi-layer anti-virus scanning. You can configure your Cisco appliance to use one or both of the licensed anti-virus scanning engines on a per mail policy basis. You could create a mail policy for executives, for example, and configure that policy to scan mail with both Sophos and McAfee engines.

Scanning messages with multiple scanning engines provides “defense in depth” by combining the benefits of both Sophos and McAfee anti-virus scanning engines. Each engine has leading anti-virus capture rates, but because each engine relies on a separate base of technology (discussed in [McAfee Anti-Virus Filtering, page 12-5](#) and [Sophos Anti-Virus Filtering, page 12-2](#)) for detecting viruses, the multi-scan approach can be even more effective. Using multiple scanning engines can lead to reduced system throughput, please contact your Cisco support representative for more information.

You cannot configure the order of virus scanning. When you enable multi-layer anti-virus scanning, the McAfee engine scans for viruses first, and the Sophos engine scans for viruses second. If the McAfee engine determines that a message is virus-free, the Sophos engine scans the message, adding a second layer of protection. If the McAfee engine determines that a message contains a virus, the Cisco appliance skips Sophos scanning and performs actions on the virus message based on settings you configured.

Sophos Anti-Virus Filtering

The Cisco appliance includes integrated virus-scanning technology from Sophos, Plc. Sophos Anti-Virus provides cross-platform anti-virus protection, detection and disinfection.

Sophos Anti-Virus provides a virus detection engine that scans files for viruses, Trojan horses, and worms. These programs come under the generic term of *malware*, meaning “malicious software.” The similarities between all types of malware allow anti-virus scanners to detect and remove not only viruses, but also all types of malicious software.

Related Topics

- [Virus Detection Engine, page 12-3](#)
- [Virus Scanning, page 12-3](#)
- [Detection Methods, page 12-3](#)
- [Virus Descriptions, page 12-4](#)
- [Sophos Alerts, page 12-4](#)
- [When a Virus is Found, page 12-4](#)

Virus Detection Engine

The Sophos virus detection engine lies at the heart of the Sophos Anti-Virus technology. It uses a proprietary architecture similar to Microsoft's COM (Component Object Model), consisting of a number of objects with well-defined interfaces. The modular filing system used by the engine is based on separate, self-contained dynamic libraries each handling a different "storage class," for example, file type. This approach allows virus scanning operations to be applied on generic data sources, irrespective of type.

Specialized technology for loading and searching data enables the engine to achieve very fast scanning speeds. Incorporated within it are:

- A full code emulator for detecting polymorphic viruses
- An on-line decompressor for scanning inside archive files
- An OLE2 engine for detecting and disinfecting macro viruses

The Cisco appliance integrates with the virus engine using SAV Interface.

Virus Scanning

In broad terms, the engine's scanning capability is managed by a powerful combination of two important components: a classifier that knows where to look, and the virus database that knows what to look for. The engine classifies the file by type rather than by relying on the extension.

The virus engine looks for viruses in the bodies and attachments of messages received by the system; an attachment's file type helps determine its scanning. For example, if a message's attached file is an executable, the engine examines the header which tells it where the executable code starts and it looks there. If the file is a Word document, the engine looks in the macro streams. If it is a MIME file, the format used for mail messaging, it looks in the place where the attachment is stored.

Detection Methods

How viruses are detected depends on their type. During the scanning process, the engine analyzes each file, identifies the type, and then applies the relevant technique(s). Underlying all methods is the basic concept of looking for certain types of instructions or certain ordering of instructions.

Related Topics

- [Pattern Matching, page 12-3](#)
- [Heuristics, page 12-4](#)
- [Emulation, page 12-4](#)

Pattern Matching

In the technique of pattern matching, the engine knows the particular sequence of code and is looking for an exact match that will identify the code as a virus. More often, the engine is looking for sequences of code that are similar, but not necessarily identical, to the known sequences of virus code. In creating the descriptions against which files are compared during scanning, Sophos virus researchers endeavor to keep the identifying code as general as possible so that – using heuristics, as explained below – the engine will find not just the original virus but also its later derivatives.

Heuristics

The virus engine can combine basic pattern matching techniques with heuristics – a technique using general rather than specific rules – to detect several viruses in the same family, even though Sophos researchers might have analyzed only one virus in that family. The technique enables a single description to be created that will catch several variants of one virus. Sophos tempers its heuristics with other methods, minimizing the incidence of false positives.

Emulation

Emulation is a technique applied by the virus engine to polymorphic viruses. Polymorphic viruses are encrypted viruses that modify themselves in an effort to hide themselves. There is no visible constant virus code and the virus encrypts itself differently each time it spreads. When it runs, it decrypts itself. The emulator in the virus detection engine is used on DOS and Windows executables, while polymorphic macro viruses are found by detection code written in Sophos's Virus Description Language.

The output of this decryption is the real virus code and it is this output that is detected by the Sophos virus detection engine after running in the emulator.

Executables that are sent to the engine for scanning are run inside the emulator, which tracks the decryption of the virus body as it is written to memory. Normally the virus entry point sits at the front end of a file and is the first thing to run. In most cases, only a small amount of the virus body has to be decrypted in order for the virus to be recognized. Most clean executables stop emulating after only a few instructions, which reduces overhead.

Because the emulator runs in a restricted area, if the code does turn out to be a virus, the virus does not infect the appliance.

Virus Descriptions

Sophos exchanges viruses with other trusted anti-virus companies every month. In addition, every month customers send thousands of suspect files directly to Sophos, about 30% of which turn out to be viruses. Each sample undergoes rigorous analysis in the highly secure virus labs to determine whether or not it is a virus. For each newly discovered virus, or group of viruses, Sophos creates a description.

Sophos Alerts

Cisco encourages customers who enable Sophos Anti-Virus scanning to subscribe to Sophos alerts on the Sophos site at <http://www.sophos.com/virusinfo/notifications/>.

Subscribing to receive alerts directly from Sophos will ensure you are apprised of the latest virus outbreaks and their available solutions.

When a Virus is Found

When a virus has been detected, Sophos Anti-Virus can repair (disinfect) the file. Sophos Anti-Virus can usually repair any file in which a virus has been found, after which the file can be used without risk. The precise action taken depends on the virus.

There can be limitations when it comes to disinfecting, because it is not always possible to return a file to its original state. Some viruses overwrite part of the executable program which cannot be reinstated. In this instance, you define how to handle messages with attachments that could not be repaired. You

configure these settings on a per-recipient basis using the Email Security Feature: the Mail Policies > Incoming or Outgoing Mail Policies pages (GUI) or the `policyconfig -> antivirus` command (CLI). For more information on configuring these settings, see [Configuring Virus Scanning Actions for Users](#), page 12-7.

McAfee Anti-Virus Filtering

The McAfee® scanning engine:

- Scans files by pattern-matching virus signatures with data from your files.
- Decrypts and runs virus code in an emulated environment.
- Applies heuristic techniques to recognize new viruses.
- Removes infectious code from files.

Related Topics

- [Pattern-Matching Virus Signatures](#), page 12-5
- [Encrypted Polymorphic Virus Detection](#), page 12-5
- [Heuristics Analysis](#), page 12-5
- [When a Virus is Found](#), page 12-6

Pattern-Matching Virus Signatures

McAfee uses anti-virus definition (DAT) files with the scanning engine to detect particular viruses, types of viruses, or other potentially unwanted software. Together, they can detect a simple virus by starting from a known place in a file, then searching for a virus signature. Often, they must search only a small part of a file to determine that the file is free from viruses.

Encrypted Polymorphic Virus Detection

Complex viruses avoid detection with signature scanning by using two popular techniques:

- **Encryption.** The data inside the virus is encrypted so that anti-virus scanners cannot see the messages or computer code of the virus. When the virus is activated, it converts itself into a working version, then executes.
- **Polymorphism.** This process is similar to encryption, except that when the virus replicates itself, it changes its appearance.

To counteract such viruses, the engine uses a technique called emulation. If the engine suspects that a file contains such a virus, the engine creates an artificial environment in which the virus can run harmlessly until it has decoded itself and its true form becomes visible. The engine can then identify the virus by scanning for a virus signature, as usual.

Heuristics Analysis

Using only virus signatures, the engine cannot detect a new virus because its signature is not yet known. Therefore the engine can use an additional technique — heuristic analysis.

Programs, documents or email messages that carry a virus often have distinctive features. They might attempt unprompted modification of files, invoke mail clients, or use other means to replicate themselves. The engine analyzes the program code to detect these kinds of computer instructions. The engine also searches for legitimate non-virus-like behavior, such as prompting the user before taking action, and thereby avoids raising false alarms.

By using these techniques, the engine can detect many new viruses.

When a Virus is Found

When a virus has been detected, McAfee can repair (disinfect) the file. McAfee can usually repair any file in which a virus has been found, after which the file can be used without risk. The precise action taken depends on the virus.

Occasionally, there can be limitations when it comes to disinfecting files because it is not always possible to return a file to its original state. Some viruses overwrite part of the executable program which cannot be reinstated. In this instance, you define how to handle messages with attachments that could not be repaired. You configure these settings on a per-recipient basis using the Email Security Feature: the Mail Policies > Incoming or Outgoing Mail Policies pages (GUI) or the `policyconfig -> antivirus` command (CLI). For more information on configuring these settings, see [Configuring Virus Scanning Actions for Users, page 12-7](#).

How to Configure the Appliance to Scan for Viruses

Table 12-1 *How to Scan Messages for Viruses*

	Do This	More Info
Step 1	Enable anti-virus scanning on the Email Security appliance.	Enabling Virus Scanning and Configuring Global Settings, page 12-7
Step 2	Define the groups of users whose messages you want to scan for viruses.	Creating a Mail Policy for a Group of Senders and Recipients, page 10-7
Step 3	(Optional) Configure how you want the virus quarantine to handle messages.	Configuring Policy, Virus, and Outbreak Quarantines, page 30-5
Step 4	Determine how you want the appliance to handle messages with viruses.	Configuring Virus Scanning Actions for Users, page 12-7
Step 5	Configure the anti-virus scanning rules for the user groups you defined.	Configuring the Anti-Virus Policies for Different Groups of Senders and Recipients, page 12-13
Step 6	(Optional) Send an email message to test the configuration.	Sending an Email to the Appliance to Test Anti-Virus Scanning, page 12-16

Related Topics

- [Enabling Virus Scanning and Configuring Global Settings, page 12-7](#)
- [Configuring Virus Scanning Actions for Users, page 12-7](#)
- [Configuring the Anti-Virus Policies for Different Groups of Senders and Recipients, page 12-13](#)
- [Notes on Anti-Virus Configurations, page 12-14](#)
- [Flow Diagram for Anti-Virus Actions, page 12-15](#)

Enabling Virus Scanning and Configuring Global Settings

You may have enabled a virus scanning engine when you ran the System Setup Wizard. Regardless, configure settings using this procedure.

**Note**

Depending on your feature keys, you can enable Sophos, McAfee, or both.

Procedure

Step 1 Navigate to the Security Services > McAfee page.

Or

Navigate to the Security Services > Sophos page.

Step 2 Click **Enable**.

**Note**

Clicking **Enable** enables the feature globally for the appliance. However, you must later enable per-recipient settings in Mail Policies.

Step 3 After reading the license agreement, scroll to the bottom of the page and click **Accept** to accept the agreement.

Step 4 Click **Edit Global Settings**.

Step 5 Choose a maximum virus scanning timeout value.

Configure a timeout value for the system to stop performing anti-virus scanning on a message. The default value is 60 seconds.

Step 6 Submit and commit your changes.

What To Do Next

Configure anti-virus settings on a per-recipient basis. See [Configuring Virus Scanning Actions for Users, page 12-7](#).

Configuring Virus Scanning Actions for Users

The virus scanning engine integrated into the Cisco appliance processes messages for viruses for incoming and outgoing mail based on policies (configuration options) you configure using the Email Security Manager feature. You enable Anti-Virus actions on a per-recipient basis using the Email Security Feature: the Mail Policies > Incoming or Outgoing Mail Policies pages (GUI) or the `policyconfig > antivirus` command (CLI).

Related Topics

- [Message Scanning Settings, page 12-8](#)
- [Message Handling Settings, page 12-8](#)
- [Configuring Settings for Message Handling Actions, page 12-9](#)

Message Scanning Settings

- Scan for Viruses Only:

Messages processed by the system are scanned for viruses. Repairs are *not* attempted for infected attachments. You can choose whether to drop attachments and deliver mail for messages that contain viruses or could not be repaired.

- Scan and Repair Viruses:

Messages processed by the system are scanned for viruses. If a virus is found in an attachment, the system will attempt to “repair” the attachment.

- Dropping Attachments

You can choose to drop infected attachments.

When infected attachments to messages have been scanned and *dropped* by the anti-virus scanning engine, the attachment is replaced with a new attachment called “Removed Attachment.” The attachment type is text/plain and contains the following:

```
This attachment contained a virus and was stripped.
```

```
Filename: filename
```

```
Content-Type: application/filetype
```

Users will always be notified if their messages were modified in any way because they were infected with a bad attachment. You can configure a secondary notification action, as well (see [Sending Notifications, page 12-11](#)). The notify action is *not* needed to inform users that a message was modified if you choose to drop infected attachments.

- X-IronPort-AV Header

All messages that are processed by the Anti-Virus scanning engine on the appliance have the header `X-IronPort-AV`: added to messages. This header provides additional information to you when debugging issues with your anti-virus configuration, particularly with messages that are considered “unscannable.” You can toggle whether the X-IronPort-AV header is included in messages that are scanned. Including this header is recommended.

Message Handling Settings

You configure the virus scanning engine to handle four distinct classes of messages that are received by a listener, with separate actions for each. [Figure 12-1](#) summarizes the actions the system performs when the virus scanning engine is enabled.

For each of the following message types, you can choose which actions are performed. The actions are described below (see [Configuring Settings for Message Handling Actions, page 12-9](#)). For example, you can configure your anti-virus settings for virus-infected messages so that the infected attachment is dropped, the subject of the email is modified, and a custom alert is sent to the message recipient.

Related Topics

- [Repaired Message Handling, page 12-9](#)
- [Encrypted Message Handling, page 12-9](#)
- [Unscannable Message Handling, page 12-9](#)
- [Virus Infected Message Handling, page 12-9](#)

Repaired Message Handling

Messages are considered *repaired* if the message was completely scanned and all viruses have been repaired or removed. These messages will be delivered as is.

Encrypted Message Handling

Messages are considered *encrypted* if the engine is unable to finish the scan due to an encrypted or protected field in the message. Messages that are marked encrypted may also be repaired.

Note the differences between the encryption detection message filter rule (see [Encryption Detection Rule, page 9-32](#)) and the virus scanning actions for “encrypted” messages. The encrypted message filter rule evaluates to “true” for any messages that are PGP or S/MIME encrypted. The encrypted rule can only detect PGP and S/MIME encrypted data. It does not detect password protected ZIP files, or Microsoft Word and Excel documents that include encrypted content. The virus scanning engine considers any message or attachment that is password protected to be “encrypted.”



Note

If you upgrade from a 3.8 or earlier version of AsyncOS and you configured Sophos Anti-Virus scanning, you must configure the Encrypted Message Handling section after you upgrade.

Unscannable Message Handling

Messages are considered *unscannable* if a scanning timeout value has been reached, or the engine becomes unavailable due to an internal error. Messages that are marked unscannable may also be repaired.

Virus Infected Message Handling

The system may be unable to drop the attachment or completely repair a message. In these cases, you can configure how the system handles messages that could still contain viruses.

The configuration options are the same for encrypted messages, unscannable messages, and virus messages.

Configuring Settings for Message Handling Actions

Related Topics

- [Action to Apply, page 12-10](#)
- [Quarantines and Anti-Virus Scanning, page 12-10](#)
- [Modify the Message Subject Header, page 12-10](#)
- [Archive Original Message, page 12-11](#)
- [Sending Notifications, page 12-11](#)
- [Add Custom Header to Message, page 12-11](#)
- [Modify message recipient, page 12-12](#)
- [Send message to alternate destination host, page 12-12](#)
- [Send custom alert notification \(to recipient only\), page 12-12](#)

Action to Apply

Choose which overall action to take on each message type for encrypted, unscannable, or virus positive messages: drop the message, deliver the message as an attachment to a new message, deliver the message as is, or send the message to the anti-virus quarantine area ([Quarantines and Anti-Virus Scanning, page 12-10](#)).

Configuring the appliance to deliver the infected messages as an attachment to a new message allows the recipient to choose how to deal with the original, infected attachment.

If you choose to deliver the message or deliver the message as an attachment to a new message, you can additionally:

- Modify message subject
 - Archive original message
 - Send generic notification
 - Add custom header to message
 - Modify message recipient
 - Send message to alternate destination host
 - Send custom alert notification (to recipient only)
- The following actions are available in the “Advanced” section of the GUI:

**Note**

These actions are not mutually exclusive; you can combine some or all of them differently within different incoming or outgoing policies for different processing needs for groups of users. See the following sections and [Notes on Anti-Virus Configurations, page 12-14](#) for more information on defining various scanning policies using these options.

**Note**

Repaired messages have only two advanced options: Add custom header and Send custom alert notification. All other message types have access to all of the advanced options.

Quarantines and Anti-Virus Scanning

When flagged for quarantine, the message continues through the rest of the email pipeline. When the message reaches the end of the pipeline, if the message has been flagged for one or more quarantines then it enters those queues. Note that if the message does not reach the end of the pipeline, it is not placed in a quarantine.

For example, a content filter can cause a message to be dropped or bounced, in which case the message will not be quarantined.

Modify the Message Subject Header

You can alter the text of identified messages by prepending or appending certain text strings to help users more easily identify and sort identified messages.

**Note**

White space is *not* ignored in the “Modify message subject” field. Add spaces after (if prepending) or before (if appending) the text you enter in this field to separate your added text from the original subject of the message. For example, add the text [WARNING: VIRUS REMOVED] with a few trailing spaces if you are prepending.

The default text is:

Table 12-2 Default Subject Line Text for Anti-Virus Subject Line Modification

Verdict	Default Text to Add to Subject
Encrypted	[WARNING: MESSAGE ENCRYPTED]
Infected	[WARNING: VIRUS DETECTED]
Repaired	[WARNING: VIRUS REMOVED]
Unscannable	[WARNING: A/V UNSCANNABLE]

Any message with multiple states causes a multi-part notification message informing users what actions the appliance performed on the message (for example, the user is notified that the message was repaired of a virus, but another part of the message was encrypted).

Archive Original Message

You can archive messages the system has identified as containing (or possibly containing) viruses to the “avarchive” directory. The format is an mbox-format log file. You *must* configure an “Anti-Virus Archive” log subscription to archive messages with viruses or messages that could not be completely scanned. For more information, see [Chapter 38, “Logging.”](#)



Note

In the GUI, you may need to click the “Advanced” link to reveal the “Archive original message” setting.

Sending Notifications

When the system has identified a message as containing viruses, you can send the default notification to the sender, the recipient, and/or additional users. When specifying additional users to notify, separate multiple addresses with commas (in both the CLI and the GUI). The default notification messages are:

Table 12-3 Default Notifications for Anti-Virus Notifications

Verdict	Notification
Repaired	The following virus(es) was detected in a mail message: <virus name(s)> Actions taken: Infected attachment dropped (or Infected attachment repaired).
Encrypted	The following message could not be fully scanned by the anti-virus engine due to encryption.
Unscannable	The following message could not be fully scanned by the anti-virus engine.
Infectious	The following unrepairable virus(es) was detected in a mail message: <virus name(s)>.

Add Custom Header to Message

You can define an additional, custom header to be added to all messages that are scanned by the anti-virus scanning engine. Click **Yes** and define the header name and text.

You can also create filters that use the `skip-viruscheck` action so that certain messages bypass virus scanning. See [Bypass Anti-Virus System Action, page 9-73](#).

Modify message recipient

You can modify the message recipient, causing the message to be delivered to a different address. Click **Yes** and enter the new recipient address.

Send message to alternate destination host


You can choose to send the notification to a different recipient or destination host for encrypted, unscannable, or virus infected messages. Click **Yes** and enter an alternate address or host.

For example, you could route suspected messages to an administrator's mailbox or a special mail server for subsequent examination. In the case of a multi-recipient message, only a single copy is sent to the alternative recipient.

Send custom alert notification (to recipient only)

You can send a custom notification to the recipient. To do so, you must first create the custom notification prior to configuring the settings. See [Understanding Text Resources, page 21-8](#) for more information.

Figure 12-1 Options for Handling Messages Scanned for Viruses

**Note**

By default, Anti-Virus scanning is enabled in the \$TRUSTED mail flow policy for public listeners, which is referenced by the WHITELIST sender group. See [Defining Access Rules for Email Senders Using Mail Flow Policies](#), page 7-8.

Configuring the Anti-Virus Policies for Different Groups of Senders and Recipients

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 Incoming or Outgoing Mail Policies. You can configure mail policies in the GUI or in the CLI using the `policyconfig > antivirus` command. 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.

Procedure

Step 1 Navigate to the Mail Policies > Incoming Mail Policies or Mail Policies > Outgoing Mail Policies page.

Step 2 Click the link for the anti-virus security service for the policy you want to configure.



Note Click the link in the default row to edit the settings for the default policy.

Step 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.

Step 4 Select an Anti-Virus scanning engine. You can select McAfee or Sophos engines.

Step 5 Configure Message Scanning settings.

See [Message Scanning Settings, page 12-8](#) for more information.

Step 6 Configure settings for Repaired, Encrypted, Unscannable, and Virus Infected messages.

See [Message Handling Settings, page 12-8](#) and [Configuring Settings for Message Handling Actions, page 12-9](#).

Step 7 Click **Submit**.

Step 8 Commit your changes.

Notes on Anti-Virus Configurations

The drop attachments flag makes a considerable difference in how anti-virus scanning works. When the system is configured to “Drop infected attachments if a virus is found and it could not be repaired,” any viral or unscannable MIME parts are removed from messages. The output from Anti-Virus scanning, then, is almost always a *clean* message. The action defined for *Unscannable Messages*, as shown in the GUI pane, rarely takes place.

In a “Scan for Viruses only” environment, these actions “clean” messages by dropping the bad message parts. Only if the RFC822 headers themselves are attacked or encounter some other problem would this result in the unscannable actions taking place. However, when Anti-Virus scanning is configured for “Scan for Viruses only” and “Drop infected attachments if a virus is found and it could not be repaired,” is *not* chosen, the unscannable actions are very likely to take place.

Table 12-4 lists some common Anti-Virus configuration options.


Table 12-4 Common Anti-Virus Configuration Options

Situation	Anti-Virus Configuration
Widespread Virus Outbreak Any viral message is simply dropped from the system with little other processing taking place.	Drop-attachments: NO Scanning: Scan-Only Cleaned messages: Deliver Unscannable messages: DROP message Encrypted messages: Send to administrator or quarantine for review. Viral messages: Drop message
Liberal Policy As many documents as possible are sent.	Drop-attachments: YES Scanning: Scan and Repair Cleaned messages: [VIRUS REMOVED] and Deliver Unscannable messages: Forward as attachment Encrypted messages: Mark and forward Viral messages: Quarantine or mark and forward.
More Conservative Policy	Drop-attachments: YES Scanning: Scan and Repair Cleaned messages: [VIRUS REMOVED] and Deliver (Archive cleaned messages for a more cautious policy.) Unscannable messages: Send notification(s), quarantine, OR drop and archive. Encrypted messages: Mark and forward OR treat as unscannable Viral messages: Archive and drop
Conservative with Review Possible virus messages are sent to a quarantine mailbox so that an administrator can review the content.	Drop-attachments: NO Scanning: Scan-Only Cleaned messages: Deliver (this action won't normally be taken) Unscannable messages: Forward as attachment, alt-src-host, or alt-rcpt-to actions. Encrypted messages: Treat as unscannable Viral messages: Forward to quarantine or administrator.

Flow Diagram for Anti-Virus Actions

Figure 12-2 on page 12-16 explains how anti-virus actions and options affect messages processed by the appliance.

Figure 12-2 Flow Diagram for Anti-Virus Actions

**Note**

If you configure multi-layer anti-virus scanning, the Cisco appliance performs virus scanning with the McAfee engine first and the Sophos engine second. It scans messages using both engines, unless the McAfee engine detects a virus. If the McAfee engine detects a virus, the Cisco appliance performs the anti-virus actions (repairing, quarantining, etc.) defined for the mail policy.

Sending an Email to the Appliance to Test Anti-Virus Scanning

Procedure

Step 1 Enable virus scanning for a mail policy.

Use the Security Services > Sophos/McAfee Anti-virus page or the `antivirusconfig` command to set global settings, and then use the Email Security Manager pages (GUI) or the `antivirus` subcommand of `policyconfig` to configure the settings for a specific mail policy.

- Step 2** Open a standard text editor, then type the following character string as *one line, with no spaces or line breaks*:

```
X5O!P%@AP[4\PZX54(P^)7CC)7}$EICAR-STANDARD-ANTIVIRUS-TEST-FILE!$H+H*
```



Note The line shown above should appear as one line in your text editor window, so be sure to maximize your text editor window and delete any line breaks. Also, be sure to type the letter O, not the number 0, in the “X5O...” that begins the test message.

If you are reading this manual on your computer, you can copy the line directly from the PDF file or HTML file and paste it into your text editor. If you copy the line, be sure to delete any extra carriage returns or spaces.

- Step 3** Save the file with the name **EICAR.COM**.

The file size will be 68 or 70 bytes.



Note This file is *not* a virus — it cannot spread or infect other files, or otherwise harm your computer. However, you should delete the file when you have finished testing your scanner to avoid alarming other users.

- Step 4** Attach the file **EICAR.COM** to an email message, and send it to the listener that will match the mail policy you configured in step 1.

Ensure that the recipient you specify in the test message will be accepted on the listener. (For more information, see [Adding Domains and Users For Which to Accept Messages, page 8-3](#).)

Note that it may be difficult to email the file if you have virus scanning software installed for outgoing mail on a gateway other than the Cisco (for example, a Microsoft Exchange server).



Note The test file always scans as unrepairable.

- Step 5** Evaluate the actions you configured for virus scanning on the listener and ensure they are enabled and working as expected.

This is most easily accomplished by performing one of the following actions:

- Configure the virus scanning settings to Scan and Repair mode or Scan only mode without dropping attachments.

Send an email with the Eicar test file as an attachment.

Confirm that the actions taken match your configuration for Virus Infected Message Handling (the settings in [Virus Infected Message Handling, page 12-9](#)).

- Configure the virus scanning settings to Scan and Repair mode or Scan only mode with dropping attachments.

Send an email with the Eicar test file as an attachment.

Confirm that the actions taken match your configuration for Repaired Message Handling (the settings in [Repaired Message Handling, page 12-9](#)).

For more information obtaining virus files for testing anti-virus scanning, see:

http://www.eicar.org/anti_virus_test_file.htm

This page provides 4 files for downloading. Note that it may be difficult to download and extract these files if you have a client-side virus scanning software installed.

Updating Virus Definitions

Related Topics

- [About Retrieving Anti-Virus Updates via HTTP](#)
- [Configuring Update Server Settings](#)
- [Monitoring and Manually Checking for Anti-Virus Updates](#)
- [Verifying Anti-Virus Files Have Updated on the Appliance](#)

About Retrieving Anti-Virus Updates via HTTP

Sophos and McAfee frequently update their virus definitions with newly-identified viruses. These updates must be passed to your appliance.

By default, the Cisco appliance is configured to check for updates every 5 minutes. For the Sophos and McAfee anti-virus engines, the server updates from a dynamic website.

The system does not timeout on updates as long as the update is actively downloading to the appliance. If the update download pauses for too long, then the download times out.

The maximum amount of time that the system waits for an update to complete before timing out is a dynamic value that is defined as 1 minute less than the anti-virus update interval (defined on Security Services > Service Updates). This configuration value aids appliances on slower connections while downloading large updates that may take longer than 10 minutes to complete.

Configuring Update Server Settings

You can configure virus update settings via the Security Services > Service Updates page. For example, you can configure how the system receives anti-virus updates and how often it checks for updates. For more information about these additional settings, see [Service Updates, page 33-17](#).

Monitoring and Manually Checking for Anti-Virus Updates

You can use the Security Services > Sophos or McAfee page or the `antivirusstatus` CLI command to verify the appliance has the latest anti-virus engine and identity files installed, and to confirm when the last update was performed.

You can also manually perform updates.

Related Topics

- [Manually Updating Anti-Virus Engines using the GUI, page 12-19](#)
- [Manually Updating Anti-Virus Engines using the CLI, page 12-19](#)

Manually Updating Anti-Virus Engines using the GUI

Procedure

-
Step 1	Navigate to the Security Services > Sophos or McAfee Anti-Virus page.
Step 2	Click Update Now in the Current McAfee/Sophos Anti-Virus Files table. The appliance checks for and downloads the latest updates.
-

Manually Updating Anti-Virus Engines using the CLI

Use the `antivirusstatus` CLI command to check the status of your virus files and `antivirusupdate` command to manually check for updates:

```
example.com> antivirusstatus
```

```
Choose the operation you want to perform:
```

- MCAFEE - Display McAfee Anti-Virus version information
- SOPHOS - Display Sophos Anti-Virus version information

```
> sophos
```

```
SAV Engine Version      3.2.07.286_4.58
  IDE Serial            0
  Last Engine Update     Base Version
  Last IDE Update        Never updated
```

```
example.com> antivirusupdate
```

```
Choose the operation you want to perform:
```

- MCAFEE - Request updates for McAfee Anti-Virus
- SOPHOS - Request updates for Sophos Anti-Virus

```
>sophos
```

```
Requesting check for new Sophos Anti-Virus updates
```

```
example.com>
```

Verifying Anti-Virus Files Have Updated on the Appliance

You can view the Updater Logs to verify whether or not the antivirus files have been successfully downloaded, extracted, or updated. Use the `tail` command to show the final entries in the Updater log subscription to ensure that virus updates were obtained.



Anti-Spam

- Overview of Anti-Spam Scanning, page 13-1
- How to Configure the Appliance to Scan Messages for Spam, page 13-2
- IronPort Anti-Spam Filtering, page 13-3
- Cisco Intelligent Multi-Scan Filtering, page 13-6
- Defining Anti-Spam Policies, page 13-7
- Protecting Appliance-Generated Messages From the Spam Filter, page 13-14
- Headers Added During Anti-Spam Scanning, page 13-14
- Reporting Incorrectly Classified Messages to Cisco Systems, page 13-15
- Determining Sender IP Address In Deployments with Incoming Relays, page 13-15
- Monitoring Rules Updates, page 13-24
- Testing Anti-Spam, page 13-25

Overview of Anti-Spam Scanning

Anti-spam processes scan email for incoming (and outgoing) mail based on the mail policies that you configure.

- One or more scanning engines scan messages through their filtering modules.
- Scanning engines assign a score to each message. The higher the score, the greater the likelihood that the message is spam.
- Based on the score, each message is categorized as one of the following:
 - Not spam
 - Unwanted marketing email from a legitimate source
 - Suspected spam
 - Positively-identified spam
- An action is taken based on the result.

Actions taken on messages positively identified as spam, suspected to be spam, or identified as unwanted marketing messages are not mutually exclusive; you can combine some or all of them differently within different incoming or outgoing policies for different processing needs for groups

of users. You can also treat positively identified spam differently from suspected spam in the same policy. For example, you may want to drop messages positively identified as spam, but quarantine suspected spam messages.

For each mail policy, you can specify thresholds for some of the categories, and determine the action to take for each category. You can assign different users to different mail policies and define different scanning engines, spam-definition thresholds, and spam-handling actions for each policy.

**Note**

For information about how and when anti-spam scanning is applied, see [Email Pipeline and Security Services, page 4-7](#).

Related Topics

- [Anti-Spam Solutions, page 13-2](#)

Anti-Spam Solutions

Your Cisco appliance offers the following anti-spam solutions:

- [IronPort Anti-Spam Filtering, page 13-3](#).
- [Cisco Intelligent Multi-Scan Filtering, page 13-6](#).

You can license and enable both these solutions on your Cisco appliance, but you can only use one in a particular mail policy. You can specify a different anti-spam solution for different groups of users.

How to Configure the Appliance to Scan Messages for Spam

	Do This	More Info
Step 1	Enable anti-spam scanning on the Email Security appliance. Note Remaining steps in this table apply to both scanning engine options.	If you have feature keys for both Cisco IronPort Anti-Spam and Intelligent Multi-Scan, you can enable both solutions on the appliance. • IronPort Anti-Spam Filtering, page 13-3 • Cisco Intelligent Multi-Scan Filtering, page 13-6
Step 2	Configure whether to quarantine spam on the local Email Security appliance or use an external quarantine on a Security Management appliance.	• Setting Up the Local Spam Quarantine, page 31-2. • Working with an External Spam Quarantine, page 42-2
Step 3	Define the groups of users whose messages you want to scan for spam.	Creating a Mail Policy for a Group of Senders and Recipients, page 10-7
Step 4	Configure the anti-spam scanning rules for the user groups you defined.	Defining Anti-Spam Policies, page 13-7
Step 5	If you want certain messages to skip Cisco Anti-Spam scanning, create message filters that use the skip-spamcheck action.	Bypass Anti-Spam System Action, page 9-72

	Do This	More Info
Step 6	(Recommended) Enable SenderBase Reputation Service scoring for each inbound mail flow policy, even if you are not rejecting connections based on SenderBase Reputation Scores.	For each inbound mail flow policy, ensure that “Use SenderBase for Flow Control” is On. See Defining Rules for Incoming Messages Using a Mail Flow Policy , page 7-15.
Step 7	If your Email Security appliance does not connect directly to external senders to receive incoming mail, but instead receives messages relayed through a mail exchange, mail transfer agent, or other machine on your network, ensure that relayed incoming messages include the original sender IP address.	Determining Sender IP Address In Deployments with Incoming Relays , page 13-15
Step 8	Prevent alert and other messages generated by your appliance from being incorrectly identified as spam.	Protecting Appliance-Generated Messages From the Spam Filter , page 13-14
Step 9	(Optional) Enable URL filtering to strengthen protection against malicious URLs in messages.	Enable URL Filtering , page 15-2
Step 10	Test your configuration.	Testing Anti-Spam , page 13-25
Step 11	(Optional) Configure settings for service updates (including anti-spam rules.)	Scanning rules for both anti-spam solutions are retrieved by default from the Cisco update servers. • Service Updates, page 33-17 • Updates Through a Proxy Server, page 33-21 • Configuring Server Settings for Downloading Upgrades and Updates, page 33-21

IronPort Anti-Spam Filtering

Related Topics

- [Evaluation Key](#), page 13-3
- [Cisco Anti-Spam: an Overview](#), page 13-4
- [Configuring IronPort Anti-Spam Scanning](#), page 13-5

Evaluation Key

Your Cisco appliance ships with a 30-day evaluation key for the Cisco Anti-Spam software. This key is not enabled until you accept the license agreement in the system setup wizard or Security Services > IronPort Anti-Spam pages (in the GUI) or the `systemsetup` or `antispsamconfig` commands (in the CLI). Once you have accepted the agreement, Cisco Anti-Spam will be enabled, by default, for the default incoming Mail Policy. An alert is also sent to the administrator address you configured (see the System Setup Wizard, [Step 2: System](#), page 3-15) noting that the Cisco Anti-Spam license will expire in 30 days. Alerts are sent 30, 15, 5, and 0 days prior to expiration. For information on enabling the feature beyond the 30-day evaluation period, contact your Cisco sales representative. You can see how much time remains on the evaluation via the System Administration > Feature Keys page or by issuing the `featurekey` command. (For more information, see [Feature Keys](#), page 33-5.)

Cisco Anti-Spam: an Overview

IronPort Anti-Spam addresses a full range of known threats including spam, phishing and zombie attacks, as well as hard-to-detect low volume, short-lived email threats such as “419” scams. In addition, IronPort Anti-Spam identifies new and evolving blended threats such as spam attacks distributing malicious content through a download URL or an executable.

To identify these threats, IronPort Anti-Spam examines the full context of a message—its content, methods of message construction, the reputation of the sender, the reputation of web sites advertised in the message, and more. IronPort Anti-Spam combines the power of email and web reputation data, leveraging the full power of the world's largest email and web traffic monitoring network — SenderBase — to detect new attacks as soon as they begin.

IronPort Anti-Spam analyzes over 100,000 message attributes across the following dimensions:

- Email reputation — *who* is sending you this message?
- Message content — *what* content is included in this message?
- Message structure — *how* was this message constructed?
- Web reputation — *where* does the call to action take you?

Analyzing multi-dimensional relationships allows the system to catch a broad range of threats while maintaining accuracy. For example, a message that has content claiming to be from a legitimate financial institution but that is sent from an IP address on a consumer broadband network or that contains a URL hosted on a “zombie” PC will be viewed as suspicious. In contrast, a message coming from a pharmaceutical company with a positive reputation will not be tagged as spam even if the message contains words closely correlated with spam.

Related Topics

- [Spam Scanning for International Regions, page 13-4](#)
- [Overview of URL Filtering, page 15-1](#)

Spam Scanning for International Regions

Cisco Anti-Spam is effective world-wide and uses locale-specific content-aware threat detection techniques. You can also optimize anti-spam scanning for a specific region using a regional rules profile.

- If you receive a large quantity of spam from a particular region outside of the US, you may want to use a regional rules profile to help you stop spam from that region.

For example, China and Taiwan receive a high percentage of spam in traditional or modern Chinese. The Chinese regional rules are optimized for this type of spam. If you receive mail primarily for mainland China, Taiwan, and Hong Kong, Cisco strongly recommends you use the Chinese regional rules profile included with the anti-spam engine.

- If your spam comes primarily from the US or from no one particular region, do not enable regional rules because doing so may reduce capture rates for other types of spam. This is because the regional rules profile optimizes the anti-spam engine for a particular region.

You can enable the regional rules profile when you configure IronPort Anti-Spam Scanning.

Related Topics

- [Configuring IronPort Anti-Spam Scanning, page 13-5](#)

Configuring IronPort Anti-Spam Scanning



Note

When IronPort Anti-Spam is enabled during system setup, it is enabled for the default incoming mail policy with the default values for the global settings.

Before You Begin

- Determine whether you will use regional scanning. See [Spam Scanning for International Regions, page 13-4](#).

Procedure

- Step 1** Select **Security Services > IronPort Anti-Spam**.
- Step 2** If you have not enabled IronPort Anti-Spam in the system setup wizard:
- Click **Enable**.
 - Scroll to the bottom of the license agreement page and click **Accept** to accept the agreement.
- Step 3** Click **Edit Global Settings**.
- Step 4** Select the check box for **Enable IronPort Anti-Spam Scanning**.
Checking this box enables the feature globally for the appliance.
- Step 5** To optimize the throughput of your appliance while still being able to scan increasingly larger messages sent by spammers, configure the thresholds for message scanning by Cisco Anti-Spam.

Option	Description
Message Scanning Thresholds	a. Enter a value for <i>Always scan messages smaller than</i>—The recommended value is 512 Kb or less. Messages smaller than the <i>always scan</i> size will be fully scanned, except in cases of “early exit.” Messages larger than this size are partially scanned if they are smaller than the <i>never scan</i> size. Cisco advises not to exceed 3 MB for the <i>always scan</i> message size. A larger value may result in decreased performance. b. Enter a value for <i>Never scan messages larger than</i>—The recommended value is 1024 Kb or less. Messages larger than this size will not be scanned by Cisco Anti-Spam and the X-IronPort-Anti-Spam-Filtered: true header will not be added to the message. Cisco advises not to exceed 10 MB for the <i>never scan</i> message size. A larger value may result in decreased performance. For messages larger than the <i>always scan</i> size or smaller than the <i>never scan</i> size, a limited and faster scan is performed. Note If the Outbreak Filters maximum message size is greater than Cisco Anti-Spam’s <i>always scan</i> message, messages smaller than the Outbreak Filters maximum size are fully scanned.

Option	Description
Timeout for Scanning Single Message	Enter the number of seconds to wait for timeout when scanning a message. Enter an integer from 1 to 120. The default value is 60 seconds.
Regional Scanning	Enable or disable regional scanning and if applicable, select a region. Enable this feature only if you receive the bulk of your email from the specified region. Because this feature optimizes the anti-spam engine for a particular region, it can reduce capture rates for other types of spam.

Step 6 Submit and commit your changes.

Cisco Intelligent Multi-Scan Filtering

Cisco Intelligent Multi-Scan incorporates multiple anti-spam scanning engines, including Cisco Anti-Spam, to provide a multi-layer anti-spam solution.

When processed by Cisco Intelligent Multi-Scan:

- A message is first scanned by third-party anti-spam engines.
- Cisco Intelligent Multi-Scan then passes the message and the verdicts of the third-party engines to Cisco Anti-Spam, which assumes responsibility for the final verdict.
- After Cisco Anti-Spam performs its scan, it returns a combined multi-scan score to AsyncOS.
- Combining the benefits of the third-party scanning engines and Cisco Anti-Spam results in more caught spam while maintaining Cisco Anti-Spam's low false positive rate.

You cannot configure the order of the scanning engines used in Cisco Intelligent Multi-Scan; Cisco Anti-Spam will always be the last to scan a message and Cisco Intelligent Multi-Scan will not skip it if a third-party engine determines that a message is spam.

Using Cisco Intelligent Multi-Scan can lead to reduced system throughput. Please contact your Cisco support representative for more information.



Note

The Intelligent Multi-Scan feature key also enables Cisco Anti-Spam on the appliance, giving you the option of enabling either Cisco Intelligent MultiScan or Cisco Anti-Spam for a mail policy.

Related Topics

- [Configuring Cisco Intelligent Multi-Scan, page 13-7](#)

Configuring Cisco Intelligent Multi-Scan

**Note**

When Cisco Intelligent Multi-Scan is enabled during system setup, it is enabled for the default incoming mail policy with the default values for the global settings.

Before You Begin

Activate the feature key for this feature. See [Feature Keys, page 33-5](#). You will see the IronPort Intelligent Multi-Scan option only if you have done so.

Procedure

-
- Step 1** Select **Security Services > IronPort Intelligent Multi-Scan**.
- Step 2** If you have not enabled Cisco Intelligent Multi-Scan in the system setup wizard:
- Click **Enable**.
 - Scroll to the bottom of the license agreement page and click **Accept** to accept the agreement.
- Step 3** Click **Edit Global Settings**.
- Step 4** Select the check box for **Enable IronPort Intelligent Multi-Scan**.
- Checking this box enables the feature globally for the appliance. However, you must still enable per-recipient settings in Mail Policies.
- Step 5** Choose a value for the *maximum message size to scan* by Cisco Intelligent Multi-Scan.
- The default value is 128 Kb. Messages larger than this size will not be scanned by Cisco Intelligent Multi-Scan.
- Step 6** Enter the number of seconds to wait for timeout when scanning a message.
- When specifying the number of seconds, enter an integer from 1 to 120. The default value is 60 seconds.
- Most users will not need to change the maximum message size to be scanned or the timeout value. That said, you may be able to optimize the throughput of your appliance by lowering the maximum message size setting.
- Step 7** Submit and commit your changes.
-

Defining Anti-Spam Policies

For each mail policy, you specify settings that determine which messages are considered spam and what action to take on those messages. You also specify which engine will scan messages that the policy applies to.

You can configure different settings for the default incoming and outgoing mail policies. If you need different anti-spam policies for different users, use multiple mail policies with different anti-spam settings. You can enable only one anti-spam solution per policy; you cannot enable both on the same policy.

Before You Begin

- Complete all steps to this point in the table in [How to Configure the Appliance to Scan Messages for Spam](#), page 13-2.
- Familiarize yourself with the following:
 - [Understanding Positive and Suspect Spam Thresholds](#), page 13-10
 - [Configuration Examples: Actions for Positively Identified versus Suspected Spam](#), page 13-11
 - [Unwanted Marketing Messages From Legitimate Sources](#), page 13-11
 - If you have enabled more than one anti-spam solution: [Enabling Different Anti-Spam Scanning Engines in Different Mail Policies: Configuration Example](#), page 13-12
 - [Headers Added During Anti-Spam Scanning](#), page 13-14
- If you will archive spam into the “Anti-Spam Archive” log, see also [Logging](#), page 38-1.
- If you will send messages to an alternate mailhost, see also [Alter Delivery Host Action](#), page 9-67.

Procedure

-
- Step 1** Navigate to the **Mail Policies > Incoming Mail Policies** page.
- Or
- Step 2** Navigate to the **Mail Policies > Outgoing Mail Policies** page.
- Step 3** Click the link under the **Anti-Spam** column for any mail policy.
- Step 4** In the **Enable Anti-Spam Scanning for This Policy** section, select the anti-spam solution you want to use for the policy.
- Options you see depend on the anti-spam scanning solution(s) that you have enabled.
- For mail policies other than the default: If you use settings from the default policy, all other options on the page are disabled.
- You can also disable anti-spam scanning altogether for this mail policy.
- Step 5** Configure settings for positively identified spam, suspected spam, and marketing messages:

Option	Description
Enable Suspected Spam Scanning	Choose an option.
Enable Marketing Email Scanning	Positively-identified spam scanning is always enabled if anti-spam scanning is enabled.
Apply This Action to Message	Choose which overall action to take on positively identified spam, suspected spam, or unwanted marketing messages: • Deliver • Drop • Bounce • Quarantine

Option	Description
(Optional) Send to Alternate Host	You can send identified messages to an alternate destination mailhost (an email server other than the ones listed in SMTP Routes or DNS). Enter an IP address or hostname. If you enter a hostname, its Mail Exchange (MX) will be queried first. If none exists, the A record on the DNS server will be used (as with SMTP Routes). Use this option if you want to redirect messages, for example to a sandbox mail server for further examination. For additional important information, see Alter Delivery Host Action, page 9-67.
Add Text to Subject	You can alter text in the Subject of identified messages by prepending or appending certain text strings to help users more easily identify and sort spam and unwanted marketing messages. Note White space is <i>not</i> ignored in this field. Add spaces after (if prepending) or before (if appending) the text you enter in this field to separate your added text from the original subject of the message. For example, if you are prepending, add the text [SPAM] with a few trailing spaces. Note “Add Text to Subject” field only accepts US-ASCII characters.
Advanced Options (for custom header and message delivery)	
(Optional) Add Custom Header	You can add a custom header to identified messages. Click Advanced and define header and value. You can use a custom header in conjunction with a content filter to perform actions such as redirecting URLs in suspected spam messages so that they pass through the Cisco Web Security proxy service. For information, see Using Custom Headers to Redirect URLs in Suspected Spam to the Cisco Web Security Proxy: Configuration Example, page 13-11.
(Optional) Send to an Alternate Envelope Recipient	You can have identified messages sent to an alternate envelope recipient address. Click Advanced and define an alternate address. For example, you could route messages identified as spam to an administrator’s mailbox for subsequent examination. In the case of a multi-recipient message, only a single copy is sent to the alternate recipient.
Archive Message	You can archive identified messages into the “Anti-Spam Archive” log. The format is an mbox-format log file.
Spam Thresholds	Use the default thresholds or enter a threshold value for positively identified spam and a value for suspected spam.

Step 6 Submit and commit your changes.

What To Do Next

If you enabled anti-spam scanning for outgoing mail, check the anti-spam settings of the relevant host access table, especially for a private listener. See [Defining Access Rules for Email Senders Using Mail Flow Policies](#), page 7-8.

Related Topics

- [How to Configure the Appliance to Scan Messages for Spam](#), page 13-2
- [Understanding Positive and Suspect Spam Thresholds](#), page 13-10
- [Configuration Examples: Actions for Positively Identified versus Suspected Spam](#), page 13-11
- [Unwanted Marketing Messages From Legitimate Sources](#), page 13-11
- [Using Custom Headers to Redirect URLs in Suspected Spam to the Cisco Web Security Proxy: Configuration Example](#), page 13-11
- [Enabling Different Anti-Spam Scanning Engines in Different Mail Policies: Configuration Example](#), page 13-12

Understanding Positive and Suspect Spam Thresholds

When evaluating messages for spam, both anti-spam scanning solutions apply thousands of rules in order to arrive at an overall spam score for the message. The score is then compared to the thresholds specified in the applicable mail policy to determine whether the message is considered spam.

For highest accuracy, the threshold for positive identification as spam is quite high by default: Messages scoring between 90 and 100 are considered to be positively identified as spam. The default threshold for suspected spam is 50.

- Messages with scores below the suspected spam threshold will be considered legitimate.
- Messages above the suspected threshold but below the positive-identification threshold will be considered to be suspected spam.

You can configure your anti-spam solution to reflect the spam tolerance levels of your organization by customizing the Positive and Suspected spam thresholds in each mail policy.

You can change the positively identified spam threshold to a value between 50 and 99. You can change the threshold for suspected spam to any value between 25 and the value you specified for positively-identified spam.

When you change the thresholds:

- Specifying a lower number (a more aggressive configuration) identifies more messages as spam and may produce more false positives. This provides a lower risk that users will see spam but a higher risk of having legitimate mail marked as spam.
- Specifying a higher number (a more conservative configuration) identifies fewer messages as spam and may deliver more spam. This provides a higher risk of users seeing spam but less risk that legitimate mail will be withheld as spam. Ideally, if set up correctly, the message subject will identify the message as likely spam and message will be delivered.

You can define separate actions to take on positively-identified and suspected spam. For example, you may want to drop “positively identified” spam but quarantine “suspected” spam.

Related Topics

- [Anti-Spam Solutions](#), page 13-2
- [Configuration Examples: Actions for Positively Identified versus Suspected Spam](#), page 13-11

Configuration Examples: Actions for Positively Identified versus Suspected Spam

Spam	Sample Actions (Aggressive)	Sample Actions (Conservative)
Positively Identified	Drop	- Deliver with “[Positive Spam]” added to the subject of messages, or - Quarantine
Suspected	Deliver with “[Suspected Spam]” added to the subject of messages	Deliver with “[Suspected Spam]” added to the subject of messages

The aggressive example tags only suspected spam messages, while dropping those messages that are positively identified. Administrators and end-users can check the subject line of incoming message for false positives, and an administrator can adjust, if necessary, the suspected spam threshold.

In the conservative example, positively identified and suspected spam is delivered with an altered subject. Users can delete suspected and positively identified spam. This method is more conservative than the first.

For a further discussion of aggressive and conservative policies in mail policies, see [Table 10-3 on page 10-11](#) in the Mail Policies chapter.

Unwanted Marketing Messages From Legitimate Sources

Both anti-spam scanning engines can distinguish between spam and unwanted marketing messages from a legitimate source. Even though marketing messages are not considered spam, your organization or end-users may not want to receive them. Like spam, you have the option to deliver, drop, quarantine, or bounce unwanted marketing messages. You also have the option to tag unwanted marketing messages by adding text to the message’s subject to identify it as marketing.

Using Custom Headers to Redirect URLs in Suspected Spam to the Cisco Web Security Proxy: Configuration Example

You can rewrite URLs in suspected spam so that when a recipient clicks a link in the message, the request is routed through the Cisco Web Security proxy service, which evaluates the safety of the site at click time and blocks access to known malicious sites.

Before You Begin

Enable the URL Filtering feature and its prerequisites. See [Setting Up URL Filtering, page 15-2](#).

Procedure

-
- Step 1** Apply a custom header to suspected spam messages:
- Select **Mail Policies > Incoming Mail Policies**.
 - Click the link in the **Anti-Spam** column for a policy such as the Default policy.

- c. In the Suspected Spam Settings section, enable suspected spam scanning.
- d. Click **Advanced** to display the Add Custom Header option.
- e. Add a custom header such as `url_redirect`.
- f. Submit and commit your changes.

Step 2 Create a content filter to redirect URLs in messages that have the custom header:

- a. Select **Mail Policies > Incoming Content Filters**.
- b. Click **Add Filter**.
- c. Name the filter `url_redirect`.
- d. Click **Add Condition**.
- e. Click **Other Header**.
- f. Enter the header name: `url_redirect`.
Make sure this exactly matches the header you created above.
- g. Select **Header exists**.
- h. Click **OK**.
- i. Click **Add Action**.
- j. Click **URL Category**.
- k. Select all categories in **Available Categories** and add them to **Selected Categories**.
- l. For Action on URL, select **Redirect to Cisco Security Proxy**.
- m. Click **OK**.

Step 3 Add the content filter to the mail policy.

- a. Select **Mail Policies > Incoming Mail Policies**.
 - b. Click the link in the **Content Filters** column for the policy that you selected earlier in this procedure.
 - a. Select **Enable Content Filters** if it is not already selected.
 - b. Select the check box to enable the **url_filtering** content filter.
 - c. Submit and commit your changes.
-

Related Topics

- [Redirecting URLs, page 14-5](#)
- [Chapter 11, “Content Filters”](#)

Enabling Different Anti-Spam Scanning Engines in Different Mail Policies: Configuration Example

When using the System Setup Wizard (or `systemsetup` command in the CLI), you are presented with option to enable either Cisco Intelligent Multi-Scan or the Cisco Anti-Spam engine. You cannot enable both during system setup, but after system setup is complete you can enable the anti-spam solution that you didn't choose, by using the Security Services menu.

After the system is set up, you can configure the anti-spam scanning solution for incoming mail policies via the Mail Policies > Incoming Mail Policies page. (Anti-spam scanning is typically disabled for outgoing mail policies.) You can even disable anti-spam scanning for a policy.

In this example, the default mail policy and the “Partners” policy are using the Cisco Anti-Spam scanning engine to quarantine positive and suspected spam.

Figure 13-1 Mail Policies - Anti-Spam Engine Per Recipient

Incoming Mail Policies

Order	Policy Name	Anti-Spam	Anti-Virus	Content Filters	Virus Outbreak Filters	Delete
1	Partners	(use default)	(use default)	(use default)	(use default)	
	Default Policy	IronPort Anti-Spam Positive: Quarantine Suspected: Quarantine	Sophos Encrypted: Deliver Unscannable: Deliver Virus Positive: Drop	Disabled	Enabled	

Key: Default Custom Disabled

To change the Partners policy to use Cisco Intelligent Multi-Scan and scan for unwanted marketing messages, click on the entry in the Anti-Spam column corresponding with the Partners row (“use default”).

Select Cisco Intelligent Multi-Scan for the scanning engine, and select Yes to enable unwanted marketing message detection. Use the default settings for unwanted marketing message detection.

Figure 13-2 shows Cisco Intelligent Multi-Scan and unwanted marketing message detection enabled in a policy.

Figure 13-2 Mail Policies - Enabling Cisco Intelligent Multi-Scan

Anti-Spam Settings

Policy: Test

Enable Anti-Spam Scanning for This Policy:

- ☐ Use Settings from Default Policy (IronPort Anti-Spam)
- ☐ Use IronPort Anti-Spam service
- ☒ Use IronPort Intelligent Multi-Scan
Spam scanning built on IronPort Anti-Spam.
- ☐ Disabled

Positively-Identified Spam Settings

Apply This Action to Message: Deliver

Send to Alternate Host (optional):

Add Text to Subject: Prepend [[SPAM]]

Advanced Optional settings for custom header and message delivery.

Suspected Spam Settings

Enable Suspected Spam Scanning: ☒ No ☒ Yes

Apply This Action to Message: Deliver

Send to Alternate Host (optional):

Add Text to Subject: Prepend [[SUSPECTED SPAM]]

Advanced Optional settings for custom header and message delivery.

Marketing Email Settings

Enable Marketing Email Scanning: ☒ No ☒ Yes

Apply This Action to Message: Deliver

Send to Alternate Host (optional):

Add Text to Subject: Prepend [[MARKETING]]

Advanced Optional settings for custom header and message delivery.

After submitting and committing the changes, the mail policy looks like this:

Figure 13-3 Mail Policies - Intelligent Multi-Scan Enabled in Policy

Incoming Mail Policies

Find Policies

Email Address:

☒ Recipient
☐ Sender

Find Policies

Policies						
Add Policy...						
Order	Policy Name	Anti-Spam	Anti-Virus	Content Filters	Virus Outbreak Filters	Delete
1	Partners	IronPort Intelligent Multi-Scan Positive: Deliver Suspected: Deliver Marketing Messages: Deliver	(use default)	(use default)	(use default)	
	Default Policy	IronPort Anti-Spam Positive: Deliver Suspected: Deliver Marketing Messages: Disabled	Not Available	Disabled	Not Available	

Key: Default Custom Disabled

Protecting Appliance-Generated Messages From the Spam Filter

Because automated email messages that are sent from the Cisco IronPort appliance (such as email alerts and scheduled reports) may contain URLs or other information that may cause them to be incorrectly identified as spam, you should do the following to ensure their delivery:

Include senders of these messages in an incoming mail policy that bypasses anti-spam scanning. See [Creating a Mail Policy for a Group of Senders and Recipients](#), page 10-7 and [Bypass Anti-Spam System Action](#), page 9-72.

Headers Added During Anti-Spam Scanning

- If either anti-spam scanning engine is enabled for a mail policy, each message that passes through that policy will have the following headers added to the message:

```
X-IronPort-Anti-Spam-Filtered: true
```

```
X-IronPort-Anti-Spam: result
```

The second header contains information that allows Cisco Support to identify the rules and engine version used to scan the message. Result information is encoded proprietary information and is not customer-decodable.

- Cisco Intelligent Multi-Scan also adds headers from the third-party anti-spam scanning engines.
- You can define additional custom headers to be added to all messages for a given mail policy that are positively identified as spam, suspected to be spam, or identified as unwanted marketing mail. See [Defining Anti-Spam Policies](#), page 13-7.

Related Topics

- [Using Custom Headers to Redirect URLs in Suspected Spam to the Cisco Web Security Proxy: Configuration Example](#), page 13-11

Reporting Incorrectly Classified Messages to Cisco Systems

Messages that appear to be incorrectly classified may be reported to Cisco for analysis. Each message is reviewed by a team of human analysts and used to enhance the accuracy and effectiveness of the product. Each message should be forwarded as an RFC 822 attachment to the following addresses:

- spam@access.ironport.com - for reporting missed spam
- ham@access.ironport.com - for reporting false-positives

Due to the volume of submissions, Cisco IronPort cannot provide individual feedback or results to customers.

For more information about reporting incorrectly classified messages, please see the Cisco Knowledge base or contact your Cisco Support provider.

Determining Sender IP Address In Deployments with Incoming Relays

If one or more mail exchange/transfer agents (MX or MTA), filtering servers, etc. stand at the edge of your network, between your Cisco appliance and the external machines that are sending incoming mail, then your appliance cannot determine the IP addresses of the sending machines. Instead, mail appears to originate from the local MX/MTA. However, IronPort Anti-Spam and Cisco Intelligent Multi-Scan (using the SenderBase Reputation Service) depend on accurate IP addresses for external senders.

The solution is to configure your appliance to work with incoming relays. You specify the names and IP addresses of all of the internal MX/MTAs connecting to the Cisco appliance, as well as the header used to store the originating IP address.

Related Topics

- [Example Environments with Incoming Relays, page 13-15](#)
- [Configuring the Appliance to Work with Incoming Relays, page 13-17](#)
- [How Incoming Relays Affect Functionality, page 13-22](#)
- [Configuring Logs to Specify Which Headers Are Used, page 13-24](#)

Example Environments with Incoming Relays

[Figure 13-4](#) shows a very basic example of an incoming relay. Mail from IP address 7.8.9.1 appears to come from IP address 10.2.3.4 because the local MX/MTA is relaying mail to the Cisco appliance.



Figure 13-4 Mail Relayed by MX/MTA – Simple

Figure 13-5 shows two other, slightly more complicated examples of how mail may be relayed inside the network and how mail may be processed by several servers within the network before it is passed to the Cisco appliance. In example A, mail from 7.8.9.1 passes through the firewall and is processed by an MX and an MTA before being delivered to the Cisco appliance. In example B, mail from 7.8.9.1 is sent to a load balancer or other type of traffic shaping appliance and is sent to any one of a range of MXs prior to being delivered to the Cisco appliance.

Figure 13-5 Mail Relayed by MX/MTA – Advanced

Configuring the Appliance to Work with Incoming Relays

Related Topics

- [Enabling the Incoming Relays Feature, page 13-17](#)
- [Adding an Incoming Relay, page 13-17](#)
- [Message Headers for Relayed Messages, page 13-18](#)

Enabling the Incoming Relays Feature

**Note**

You should only enable the incoming relays feature if a local MX/MTA relays mail to your Cisco appliance.

Procedure

-
- Step 1** Select **Network > Incoming Relays**.
- Step 2** Click **Enable**.
- Step 3** Commit your changes.
-

Adding an Incoming Relay

Add incoming relays to identify:

- Each machine on your network that will relay incoming messages to your Email Security appliance, and
- The header that will label the IP address of the original external sender.

Before You Begin

For information needed to complete these prerequisites, see [Message Headers for Relayed Messages, page 13-18](#).

- Determine whether you will use custom or received headers to identify the IP address of the original external sender.
- If you will use custom headers:
 - Determine the exact header that will label the originating IP address of relayed messages.
 - For each MX, MTA, or other machine that connects to original external senders, set up that machine to add the header name and the IP address of the original external sender to incoming messages.

Procedure

-
- Step 1** Select **Network > Incoming Relays**.
- Step 2** Click **Add Relay**.
- Step 3** Enter a name for this relay.

Step 4 Enter the IP address of the MTA, MX, or other machine that connects to the Email Security appliance to relay incoming messages.

You can use IPv4 or IPv6 addresses, standard CIDR format, or an IP address range. For example, if you have several MTAs at the edge of your network receiving email, you might want to enter a range of IP addresses to include all of your MTAs, such as 10.2.3.1/8 or 10.2.3.1-10.

For IPv6 addresses, AsyncOS supports the following formats:

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

Step 5 Specify the header that will identify the IP address of the original external sender.

When entering a header, you do not need to enter the trailing colon.

a. Select the header type:

Choose custom headers (recommended) or Received headers.

b. For custom headers:

Enter the header name that you configured the relaying machine to add to relayed messages.

For example:

`SenderIP`

or

`X-CustomHeader`

c. 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.

Step 6 Submit and commit your changes.

What To Do Next

Consider doing the following:

- Add the relaying machine to a sender group with a mail flow policy that has unlimited messages for DHAP. For an explanation, see [Incoming Relays and Directory Harvest Attack Prevention, page 13-22](#).
- To facilitate tracking and troubleshooting, configure the appliance logs to show which header is used. See [Configuring Logs to Specify Which Headers Are Used, page 13-24](#).

Related Topics

- [How to Configure the Appliance to Scan Messages for Spam, page 13-2](#)

Message Headers for Relayed Messages

You will configure your appliance to use one of the following types of header to identify the original sender of a relayed message:

- [Custom Header, page 13-19](#)

- [Received Header, page 13-19](#)

Custom Header


Using custom headers is the recommended method of identifying original senders. The machine connecting to the original sender needs to add this custom header. The value of the header is expected to be the IP address of the external sending machine. For example:

```
SenderIP: 7.8.9.1
```

```
X-CustomHeader: 7.8.9.1
```

If your local MX/MTA can receive mail from a variable number of hops, inserting a custom header is the only way to enable the Incoming Relays feature. For example, in [Figure 13-6](#) both path C and D lead to IP address 10.2.3.5; however, path C has two hops and path D has one. Because the number of hops can vary in this situation, you must use a custom header in order to have Incoming Relays configured correctly.

Figure 13-6 Mail Relayed by MX/MTA — Variable Number of Hops



Related Topics

- [Adding an Incoming Relay, page 13-17](#)

Received Header

If configuring the MX/MTAs to include a custom header containing the sending IP address is not an option, you can configure the incoming relays feature to attempt to determine the sending IP address by examining the “Received:” headers in the message. Using the “Received:” header will only work if the number of network “hops” will always be constant for an IP address. In other words, the machine at the first hop (10.2.3.5 in [Figure 13-5](#)) should always be the same number of hops away from the edge of your

network. If incoming mail can take different paths (resulting in a different number of hops, as described in [Figure 13-6](#)) to the machine connecting to your Cisco appliance, you must use a custom header (see [Custom Header](#), [page 13-19](#)).

Specify a parsing character or string and the number of network hops (or Received: headers) back to look. A hop is basically the message traveling from one machine to another (being received by the Cisco appliance does not count as a hop. See [Configuring Logs to Specify Which Headers Are Used](#), [page 13-24](#) for more information). AsyncOS looks for the first IP address following the first occurrence of the parsing character or string in the Received: header corresponding to the number of specified hops. For example, if you specify two hops, the second Received: header, working backward from the Cisco appliance is parsed. If neither the parsing character nor a valid IP address is found, the Cisco appliance uses the real IP address of the connecting machine.

For the following example mail headers, if you specify an opening square bracket ([) and two hops, the IP address of the external machine is 7.8.9.1. However, if you specify an closing parenthesis ()) as the parsing character, a valid IP address will not be found. In this case, the Incoming Relays feature is treated as disabled, and the IP of the connecting machine is used (10.2.3.5).

In the example in [Figure 13-5](#) the incoming relays are:

- Path A — 10.2.3.5 (with 2 hops when using received headers) and
- Path B — 10.2.6.1 (with 2 hops when using received headers)

[Table 13-1](#) shows example email headers for a message as it moves through several hops on its way to the Cisco appliance as in [Figure 13-5](#). This example shows extraneous headers (ignored by your Cisco appliance) which are present once the message has arrived in the recipient's inbox. The number of hops to specify would be two. [Table 13-2](#) shows the headers for the same email message, without the extraneous headers

Table 13-1 A Series of Received: Headers (Path A Example 1)

1	Microsoft Mail Internet Headers Version 2.0 Received: from smemail.rand.org ([10.2.2.7]) by smmail5.customerdomain.org with Microsoft SMTPSVC(5.0.2195.6713); Received: from ironport.customerdomain.org ([10.2.3.6]) by smemail.customerdomain.org with Microsoft SMTPSVC(5.0.2195.6713);
2	Received: from mta.customerdomain.org ([10.2.3.5]) by ironport.customerdomain.org with ESMTP; 21 Sep 2005 13:46:07 -0700
3	Received: from mx.customerdomain.org (mx.customerdomain.org) [10.2.3.4] by mta.customerdomain.org (8.12.11/8.12.11) with ESMTP id j8LKKWu1008155 for <joefoo@customerdomain.org>

Table 13-1 A Series of Received: Headers (Path A Example 1) (continued)

4	Received: from sending-machine.spamham.com (sending-machine.spamham.com [7.8.9.1]) by mx.customerdomain.org (Postfix) with ESMTTP id 4F3DA15AC22 for <joefoo@customerdomain.org>
5	Received: from linux1.thespammer.com (HELO linux1.thespammer.com) ([10.1.1.89]) by sending-machine.spamham.com with ESMTTP; Received: from exchange1.thespammer.com ([10.1.1.111]) by linux1.thespammer.com with Microsoft SMTPSVC(6.0.3790.1830); Subject: Would like a bigger paycheck? Date: Wed, 21 Sep 2005 13:46:07 -0700 From: "A. Sender" <asend@otherdomain.com> To: <joefoo@customerdomain.org>

Notes for [Table 13-1](#):

- The Cisco appliance ignores these headers.
- The Cisco appliance receives the message (not counted as a hop).
- First hop (and incoming relay).
- Second hop. This is the sending MTA. The IP address is 7.8.9.1.
- The Cisco appliance ignores these Microsoft Exchange headers.

Table 13-2 A Series of Received: Headers (Path A Example 2)

1	Received: from mta.customerdomain.org ([10.2.3.5]) by ironport.customerdomain.org with ESMTTP; 21 Sep 2005 13:46:07 -0700
2	Received: from mx.customerdomain.org (mx.customerdomain.org) [10.2.3.4]) by mta.customerdomain.org (8.12.11/8.12.11) with ESMTTP id j8LKkWu1008155 for <joefoo@customerdomain.org>;
3	Received: from sending-machine.spamham.com (sending-machine.spamham.com [7.8.9.1]) by mx.customerdomain.org (Postfix) with ESMTTP id 4F3DA15AC22 for <joefoo@customerdomain.org>;

[Figure 13-7](#) shows the incoming relay for path A (above) as configured in the Add Relay page in the GUI:

Figure 13-7 A Configured Incoming Relay with Received Header**Add Relay**

Incoming Relay	
Name: ?	IncomingRelayOne
IP Address: ?	10.2.3.5
Header:	☐ Specify a custom header
	☒ Parse the "Received" header
	Begin parsing after: ?
	Hop: ?

Related Topics

- [Adding an Incoming Relay, page 13-17](#)

How Incoming Relays Affect Functionality

- [Incoming Relays and Filters, page 13-22](#)
- [Incoming Relays, HAT, SBRS, and Sender Groups, page 13-22](#)
- [Incoming Relays and Directory Harvest Attack Prevention, page 13-22](#)
- [Incoming Relays and Trace, page 13-23](#)
- [Incoming Relays and Email Security Monitor \(Reporting\), page 13-23](#)
- [Incoming Relays and Message Tracking, page 13-23](#)
- [Incoming Relays and Logging, page 13-23](#)

Incoming Relays and Filters

The Incoming Relays feature provides the various SenderBase Reputation Service related filter rules (`reputation`, `no-reputation`) with the correct SenderBase Reputation score.

Incoming Relays, HAT, SBRS, and Sender Groups

HAT policy groups do not currently use information from Incoming Relays. However, because the Incoming Relays feature does supply the SenderBase Reputation score, you can simulate HAT policy group functionality via message filters and the `$reputation` variable.

Incoming Relays and Directory Harvest Attack Prevention

If a remote host attempts a directory harvest attack by sending messages to the MX or MTA serving as an incoming relay on your network, the appliance drops the connection from the incoming relay if the relay is assigned to a sender group with a mail flow policy with Directory Harvest Attack Prevention (DHAP) enabled. This prevents all messages from the relay, including legitimate messages, from reaching the Email Security appliance. The appliance does not have the opportunity to recognize the remote host as the attacker and the MX or MTA that's acting as the incoming relay continues to receive

mail from the attacking host. To work around this issue and continue receiving messages from the incoming relay, add the relay to a sender group with a mail flow policy that has unlimited messages for DHAP.

Incoming Relays and Trace

Trace returns the Incoming Relay's SenderBase Reputation Score in its results instead of the reputation score for the source IP address.

Incoming Relays and Email Security Monitor (Reporting)

When using Incoming Relays:

- Email Security Monitor reports include data for both the external IP and the MX/MTA. For example, if an external machine (IP 7.8.9.1) sent 5 emails through the internal MX/MTA (IP 10.2.3.4), Mail Flow Summary will show 5 messages coming from IP 7.8.9.1 and 5 more coming from the internal relay MX/MTA (IP 10.2.3.5).
- The SenderBase Reputation score is not reported correctly in the Email Security Monitor reports. Also, sender groups may not be resolved correctly.

Incoming Relays and Message Tracking

When using Incoming Relays, the Message Tracking Details page displays the relay's IP address and the relay's SenderBase Reputation Score for a message instead of the IP address and reputation score of the original external sender.

Incoming Relays and Logging

In the following log example, the SenderBase Reputation score for the sender is reported initially on line 1. Later, once the Incoming Relay is processed, the correct SenderBase Reputation score is reported on line 5.

1	Fri Apr 28 17:07:29 2006 Info: ICID 210158 ACCEPT SG UNKNOWNLIST match nx.domain SBRS rfc1918
2	Fri Apr 28 17:07:29 2006 Info: Start MID 201434 ICID 210158
3	Fri Apr 28 17:07:29 2006 Info: MID 201434 ICID 210158 From: <joe@sender.com>
4	Fri Apr 28 17:07:29 2006 Info: MID 201434 ICID 210158 RID 0 To: <mary@example.com>
5	Fri Apr 28 17:07:29 2006 Info: MID 201434 IncomingRelay(senderdotcom): Header Received found, IP 192.192.108.1 being used, SBRS 6.8
6	Fri Apr 28 17:07:29 2006 Info: MID 201434 Message-ID '<7.0.1.0.2.20060428170643.0451be40@sender.com>'
7	Fri Apr 28 17:07:29 2006 Info: MID 201434 Subject 'That report...'
8	Fri Apr 28 17:07:29 2006 Info: MID 201434 ready 2367 bytes from <joe@sender.com>
9	Fri Apr 28 17:07:29 2006 Info: MID 201434 matched all recipients for per-recipient policy DEFAULT in the inbound table
10	Fri Apr 28 17:07:34 2006 Info: ICID 210158 close
11	Fri Apr 28 17:07:35 2006 Info: MID 201434 using engine: CASE spam negative

12	Fri Apr 28 17:07:35 2006 Info: MID 201434 antivirus negative
13	Fri Apr 28 17:07:35 2006 Info: MID 201434 queued for delivery

Incoming Relays and Mail Logs

The following example shows a typical log entry containing Incoming Relay information:

```
Wed Aug 17 11:20:41 2005 Info: MID 58298 IncomingRelay(myrelay): Header Received found,
IP 192.168.230.120 being used
```

Configuring Logs to Specify Which Headers Are Used

Your Cisco appliance only examines the headers that were present when the message was received. So, additional headers added locally (such as Microsoft Exchange headers, etc.) or when the message is received by the Cisco appliance are not processed. One way to help determine which headers are used is to configure AsyncOS logging to include the headers you use.

To configure logging settings for headers, see [Configuring Global Settings for Logging, page 38-40](#).

Monitoring Rules Updates

Once you have accepted the license agreement, you can view the most recent Cisco Anti-Spam and Cisco Intelligent Multi-Scan rules updates.

Procedure

- Step 1** Select Security Services > IronPort Anti-Spam.
or
- Step 2** Select Security Services > IronPort Intelligent Multi-Scan.
- Step 3** Look at the **Rule Updates** section and:

To	More Information
See the most recent update for each component	If an update has not occurred, or a server has not been configured, “Never Updated” is displayed.
See if an update is available	—
Update rules if updates are available	Click Update Now .

Related Topics

- [Service Updates, page 33-17](#)
- [Updates Through a Proxy Server, page 33-21](#)
- [Configuring Server Settings for Downloading Upgrades and Updates, page 33-21](#)

Testing Anti-Spam

To	Do This	More Information
Test your configuration.	Test your configuration using the <code>X-advertisement: spam</code> header. For testing purposes, Cisco Anti-Spam considers any message with an X-header formatted as <code>X-Advertisement: spam to be spam</code>.	The test message you send with this header is flagged by Cisco Anti-Spam, and you can confirm that the actions you configured for the mail policy (Defining Anti-Spam Policies, page 13-7) are performed. Use this header with one of the following: • Use SMTP commands to send a test message with this header. See Sending an Email to the Appliance to Test Cisco Anti-Spam, page 13-25. • Use the <code>trace</code> command and include this header. See Debugging Mail Flow Using Test Messages: Trace, page 40-1.
Evaluate Anti-Spam engine efficacy.	Evaluate the product using a live mail stream directly from the Internet.	For a list of ineffective evaluation approaches that you should avoid, see Ways Not to Test Anti-Spam Efficacy, page 13-26 .

Related Topics

- [Sending an Email to the Appliance to Test Cisco Anti-Spam, page 13-25](#)
- [Ways Not to Test Anti-Spam Efficacy, page 13-26](#)

Sending an Email to the Appliance to Test Cisco Anti-Spam

Before You Begin

- Understand how to use Telnet with the appliance. See [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#)
- Review the example in [Testing Anti-Spam Configuration: Example Using SMTP, page 13-26](#).

Procedure

-
- Step 1** Enable Cisco Anti-Spam on a mail policy.
- Step 2** Send a test email that includes the following header to a user in that mail policy:
- ```
X-Advertisement: spam
```
- Use SMTP commands with Telnet to send this message to an address to which you have access.
- Step 3** Check the mailbox of the test account and confirm that the test message was correctly delivered based upon the actions you configured for the mail policy.
- For example:
- Was the subject line altered?
  - Was your additional custom header added?
  - Was the message delivered to an alternate address?
  - Was the message dropped?

**Related Topics**

- [Testing Anti-Spam Configuration: Example Using SMTP, page 13-26](#)

## Testing Anti-Spam Configuration: Example Using SMTP

For this example, the mail policy must be configured to receive messages for the test address and the HAT must accept the test connection.

```
telnet IP_address_of_IronPort_Appliance_with_IronPort_Anti-Spam port

220 hostname ESMTP

helo example.com

250 hostname

mail from: <test@example.com>

250 sender <test@example.com> ok

rcpt to: <test@address>

250 recipient <test@address> ok

data

354 go ahead

Subject: Spam Message Test

X-Advertisement: spam

spam test

.

250 Message MID accepted

221 hostname

quit
```

## Ways Not to Test Anti-Spam Efficacy

Because IronPort AntiSpam and Cisco Intelligent Multi-Scan rules are added quickly to prevent active spam attacks and quickly expire once attacks have passed, you should not test efficacy using any of the following methods:

- Evaluating using resent or forwarded mail or cut-and-pasted spam messages.  
Mail lacking the proper headers, connecting IP, signatures, etc. will result in inaccurate scores.
- Testing “hard spam” only.

Removing the “easy spam” using SBRS, blacklists, message filters, etc. will result in a lower overall catch rate percentage.

- Resending spam caught by another anti-spam vendor.
- Testing older messages.

The scanning engine adds and removes rules rapidly based on current threats. Testing using old messages will therefore lead to inaccurate test results.



## Outbreak Filters

---

- [Overview of Outbreak Filters, page 14-1](#)
- [How Outbreak Filters Work, page 14-2](#)
- [How the Outbreak Filters Feature Works, page 14-8](#)
- [Managing Outbreak Filters, page 14-11](#)
- [Monitoring Outbreak Filters, page 14-23](#)
- [Troubleshooting The Outbreak Filters Feature, page 14-24](#)

### Overview of Outbreak Filters

Outbreak Filters protects your network from large-scale virus outbreaks and smaller, non-viral attacks, such as phishing scams and malware distribution, as they occur. Unlike most anti-malware security software, which cannot detect new outbreaks until data is collected and a software update is published, Cisco gathers data on outbreaks as they spread and sends updated information to your Email Security appliance in real-time to prevent these messages from reaching your users.

Cisco uses global traffic patterns to develop rules that determine if an incoming message is safe or part of an outbreak. Messages that may be part of an outbreak are quarantined until they're determined to be safe based on updated outbreak information from Cisco or new anti-virus definitions are published by Sophos and McAfee.

Messages used in small-scale, non-viral attacks use a legitimate-looking design, the recipient's information, and custom URLs that point to phishing and malware websites that have been online only for a short period of time and are unknown to web security services. Outbreak Filters analyzes a message's content and searches for URL links to detect this type of non-viral attack. Outbreak Filters can rewrite URLs to redirect traffic to potentially harmful websites through a web security proxy, which either warns users that the website they are attempting to access may be malicious or blocks the website completely.



# How Outbreak Filters Work

## Related Topics

- [Delaying, Redirecting, and Modifying Messages, page 14-2](#)
- [Threat Categories, page 14-2](#)
- [Cisco Security Intelligence Operations, page 14-3](#)
- [Context Adaptive Scanning Engine, page 14-4](#)
- [Delaying Messages, page 14-4](#)
- [Redirecting URLs, page 14-5](#)
- [Modifying Messages, page 14-6](#)
- [Types of Rules: Adaptive and Outbreak, page 14-6](#)
- [Outbreaks, page 14-7](#)
- [Threat Levels, page 14-7](#)

## Delaying, Redirecting, and Modifying Messages

The Outbreak Filters feature uses three tactics to protect your users from outbreaks:

- **Delay.** Outbreak Filters quarantines messages that may be part of a virus outbreak or non-viral attack. While quarantined, the appliance receives updated outbreak information and rescans the message to confirm whether it's part of an attack.
- **Redirect.** Outbreak Filters rewrites the URLs in non-viral attack messages to redirect the recipient through the Cisco web security proxy if they attempt to access any of the linked websites. The proxy displays a splash screen that warns the user that the website may contain malware, if the website is still operational, or displays an error message if the website has been taken offline. See [Redirecting URLs, page 14-5](#) for more information on redirecting URLs.
- **Modify.** In addition to rewriting URLs in non-viral threat messages, Outbreak Filters can modify a message's subject and add a disclaimer above the message body to warn users about the message's content. See [Modifying Messages, page 14-6](#) for more information.

## Threat Categories

The Outbreak Filters feature provides protection from two categories of message-based outbreaks: *virus outbreaks*, which are messages with never-before-seen viruses in their attachments, and *non-viral threats*, which includes phishing attempts, scams, and malware distribution through links to an external website.

By default, the Outbreak Filters feature scans your incoming and outgoing messages for possible viruses during an outbreak. You can enable scanning for non-viral threats in addition to virus outbreaks if you enable anti-spam scanning on the appliance.



### Note

Your appliance needs a feature key for Anti-Spam or Intelligent Multi-Scan in order for Outbreak Filters to scan for non-viral threats.

**Related Topics**

- [Virus Outbreaks, page 14-3](#)
- [Phishing, Malware Distribution, and Other Non-Viral Threats, page 14-3](#)

## Virus Outbreaks

The Outbreak Filters feature provides you with a head start when battling virus outbreaks. An outbreak occurs when messages with attachments containing never-before-seen viruses or variants of existing viruses spread quickly through private networks and the Internet. As these new viruses or variants hit the Internet, the most critical period is the window of time between when the virus is released and when the anti-virus vendors release an updated virus definition. Having advanced notice — even a few hours — is vital to curbing the spread of the malware or virus. During that vulnerability window, the newly-found virus can propagate globally, bringing email infrastructure to a halt.

## Phishing, Malware Distribution, and Other Non-Viral Threats

Messages containing non-viral threats are designed to look like a message from a legitimate sources and often sent out to a small number of recipients. These messages may have one or more of the following characteristics in order to appear trustworthy:

- The recipient's contact information.
- HTML content designed to mimic emails from legitimate sources, such as social networks and online retailers.
- URLs pointing to websites that have new IP addresses and are online only for a short time, which means that email and web security services do not have enough information on the website to determine if it is malicious.
- URLs pointing to URL shortening services.

All of these characteristics make these messages more difficult to detect as spam. The Outbreak Filters feature provides a multi-layer defense from these non-viral threats to prevent your users from downloading malware or providing personal information to suspicious new websites.

If CASE discovers URLs in the message, it compares the message to existing Outbreak Rules to determine if the message is part of a small-scale non-viral outbreak and then assigns a threat level. Depending on the threat level, the Email Security appliance delays delivery to the recipient until more threat data can be gathered and rewrites the URLs in the message to redirect the recipient to the Cisco web security proxy if they attempt to access the website. The proxy displays a splash page warning the user that the website may contain malware.

## Cisco Security Intelligence Operations

Cisco Security Intelligence Operations (SIO) is a security ecosystem that connects global threat information, reputation-based services, and sophisticated analysis to Cisco security appliances to provide stronger protection with faster response times.

SIO consists of three components:

- SenderBase. The world's largest threat monitoring network and vulnerability database.
- Threat Operations Center (TOC). A global team of security analysts and automated systems that extract actionable intelligence gathered by SenderBase.
- Dynamic Update. Real-time updates automatically delivered to appliances as outbreaks occur.

SIO compares real-time data from the global SenderBase network to common traffic patterns to identify anomalies that are proven predictors of an outbreak. TOC reviews the data and issues a threat level of the possible outbreak. Cisco Email Security appliances download updated threat levels and Outbreak Rules and use them to scan incoming and outgoing messages, as well as messages already in the Outbreak quarantine.

Information about current virus outbreaks can be found on SenderBase's website here:

<http://www.senderbase.org/>

The SIO website provides a list of current non-viral threats, including spam, phishing, and malware distribution attempts:

<http://tools.cisco.com/security/center/home.x>

## Context Adaptive Scanning Engine

Outbreak Filters are powered by Cisco's unique Context Adaptive Scanning Engine (CASE). CASE leverages over 100,000 adaptive message attributes tuned automatically and on a regular basis, based on real-time analysis of messaging threats.

For virus outbreaks, CASE analyzes the message content, context and structure to accurately determine likely Adaptive Rule triggers. CASE combines Adaptive Rules and the real-time Outbreak Rules published by SIO to evaluate every message and assign a unique threat level.

To detect non-viral threats, CASE scans messages for URLs and uses Outbreak Rules from SIO to evaluate a message's threat level if one or more URLs are found.

Based on the message's threat level, CASE recommends a period of time to quarantine the message to prevent an outbreak. CASE also determines the rescan intervals so it can reevaluate the message based on updated Outbreak Rules from SIO. The higher the threat level, the more often it rescans the message while it is quarantined.

CASE also rescans messages when they're released from the quarantine. A message can be quarantined again if CASE determines that it is spam or contains a virus upon rescan.

For more information about CASE, see [Cisco Anti-Spam: an Overview, page 13-4](#).

## Delaying Messages

The period between when an outbreak or email attack occurs and when software vendors release updated rules is when your network and your users are the most vulnerable. A modern virus can propagate globally and a malicious website can deliver malware or collect your users' sensitive information during this period. Outbreak Filters protects your users and network by quarantining suspect messages for a limited period of time, giving Cisco and other vendors time to investigate the new outbreak.

When a virus outbreak occurs, suspicious messages with attachments are quarantined until updated Outbreak Rules and new anti-virus signatures prove the email's attachment is clean or a virus.

Small scale, non-viral threats contain URLs to malicious websites that may be online for a short period of time in order to evade detection by web security services or through URL shortening services in order to circumvent web security by putting a trustworthy website in the middle. By quarantining messages containing URLs that meet your threat level threshold, not only does CASE have the opportunity to reevaluate the message's content based on updated Outbreak Rules from SIO, but the messages can remain in the quarantine long enough that the linked website may go offline or be blocked by a web security solution.

See [Dynamic Quarantine, page 14-10](#) for more information on how Outbreak Filters quarantine suspicious messages.

## Redirecting URLs

When CASE scans a message at the Outbreak Filters stage, it searches for URLs in the message body in addition to other suspicious content. CASE uses published Outbreak Rules to evaluate whether the message is a threat and then scores the message with the appropriate threat level. Depending on the threat level, Outbreak Filters protects the recipient by rewriting all the URLs to redirect the recipient to the Cisco web security proxy, except for URLs pointing to bypassed domains, and delaying the delivery of the message in order for TOC to learn more about the website if it appears to be part of a larger outbreak. See [URL Rewriting and Bypassing Domains, page 14-20](#) for more information on bypassing URLs for trusted domains.

After the Email Security appliance releases and delivers the message, any attempt by the recipient to access the website is redirected through the Cisco web security proxy. This is an external proxy hosted by Cisco that displays a splash screen that warns the user that the website may be dangerous, if the website is still operational. If the website has been taken offline, the splash screen displays an error message.

If the recipient decides to click the message's URLs, the Cisco web security proxy displays a splash screen in the user's web browser to warn the user about the content of the message. [Figure 14-1](#) shows an example of the splash screen warning. The recipient can either click **Ignore this warning** to continue on to the website or **Exit** to leave and safely close the browser window.

**Figure 14-1** Cisco Security Splash Screen Warning



### The requested web page may be dangerous

Cisco Email and Web Security protects your organization's network from malicious software. Malware is designed to look like a legitimate email or website which accesses your computer, hides itself in your system, and damages files. Your email administrator has configured this prevention system to ensure against such damage.

Exit

[Ignore this warning](#)

The only way to access the Cisco web security proxy is through a rewritten URL in a message. You cannot access the proxy by typing a URL in your web browser.



#### Note

You can customize the appearance of this splash screen and display your organization's branding such as company logo, contact information, and so on. See [Customizing the Appearance of End User Notification Page, page 15-6](#).

**Tip**

To redirect all URLs in suspected spam messages to the Cisco Web Security proxy service, see [Using Custom Headers to Redirect URLs in Suspected Spam to the Cisco Web Security Proxy: Configuration Example, page 13-11](#).

## Modifying Messages

The Outbreak Filters feature modifies the message body of a non-viral threat message not only to rewrite the URLs but to alert the user that the message is a suspected threat. The Outbreak Filters feature can modify the subject header and add a disclaimer about the message's content above the message body. See [Message Modification, page 14-18](#) for more information.

The threat disclaimer is created using the Disclaimer template through the Mail Policies > Text Resources page. See [Overview of Text Resource Management, page 21-9](#) for more information.

## Types of Rules: Adaptive and Outbreak

Two types of rules are used by Outbreak Filters to detect potential outbreaks: Adaptive and Outbreak. The Outbreak Filters feature uses these two rule sets to provide the highest efficacy and the most focused set of criteria for threat detection to ensure that filters can be laser focused on a particular outbreak. The Outbreak Filters rules and actions are visible to the administrator, not hidden away behind the scenes, providing instant access to quarantined messages and the reason why they were quarantined.

### Related Topics

- [Outbreak Rules, page 14-6](#)
- [Adaptive Rules, page 14-7](#)

## Outbreak Rules

Outbreak Rules are generated by the Cisco Threat Operations Center (TOC), which is a part of the Cisco Security Intelligence Operations, and focus on the message as a whole, rather than just attachment filetypes. Outbreak Rules use SenderBase data (real time and historical traffic data) and any combination of message parameters such as attachment file type, file name keywords, or anti-virus engine update to recognize and prevent outbreaks in real time. Outbreak Rules are given a unique ID used to refer to the rule in various places in the GUI (such as the Outbreak quarantine).

Real-time data from the global SenderBase network is then compared to this baseline, identifying anomalies that are proven predictors of an outbreak. The TOC reviews the data and issues a threat indicator or Threat Level. The Threat Level is a numeric value between 0 (no threat) and 5 (extremely risky), and measures the likelihood that a message is a threat for which no other gateway defense is widely deployed by Cisco customers (for more information, see [Threat Levels, page 14-7](#)). Threat Levels are published as Outbreak Rules by the TOC.

Some example characteristics that can be combined in Outbreak Rules include:

- File Type, File Type & Size, File Type & File Name Keyword, etc.
- File Name Keyword & File Size
- File Name Keyword
- Message URL

- File Name & Sophos IDE

## Adaptive Rules

Adaptive Rules are a set of rules within CASE that accurately compare message attributes to attributes of known virus outbreak messages. These rules have been created after studying known threat messages and known good messages within an extensive virus corpus. Adaptive Rules are updated often as the corpus is evaluated. They complement existing Outbreak Rules to detect outbreak messages at all times. While Outbreak Rules take effect when a possible outbreak is occurring, Adaptive Rules (once enabled) are “always on,” catching outbreak messages locally before the full anomaly has formed on a global basis. Additionally, Adaptive Rules continuously respond to small and subtle changes in email traffic and structure, providing updated protection to customers.

## Outbreaks

A Outbreak Filter rule is basically a Threat Level (e.g. 4) associated with a set of characteristics for an email message and attachment — things such as file size, file type, file name, message content, and so on. For example, assume the Cisco SIO notices an increase in the occurrences of a suspicious email message carrying a .exe attachment that is 143 kilobytes in size, and whose file name includes a specific keyword (“hello” for example). An Outbreak Rule is published increasing the Threat Level for messages matching this criteria. Your appliance checks for and downloads newly published Outbreak and Adaptive Rules every 5 minutes by default (see [Updating Outbreak Filter Rules, page 14-15](#)). Adaptive Rules are updated less frequently than Outbreak Rules. On the appliance, you set a threshold for quarantining suspicious messages. If the Threat Level for a message equals or exceeds the quarantine threshold, the message is sent to the *Outbreak* quarantine area. You can also set up a threshold for modifying non-viral threat messages to rewrite any URLs found in suspicious messages or add a notification at the top of message body.

## Threat Levels

[Table 14-1 on page 14-7](#) provides a basic set of guidelines or definitions for each of the various levels.

**Table 14-1**      **Threat Level Definitions**

| Level | Risk       | Meaning                                                                                                                                |
|-------|------------|----------------------------------------------------------------------------------------------------------------------------------------|
| 0     | None       | There is no risk that the message is a threat.                                                                                         |
| 1     | Low        | The risk that the message is a threat is low.                                                                                          |
| 2     | Low/Medium | The risk that the message is a threat is low to medium. It is a “suspected” threat.                                                    |
| 3     | Medium     | Either the message is part of a confirmed outbreak or there is a medium to large risk of its content being a threat.                   |
| 4     | High       | Either the message is confirmed to be part of a large scale outbreak or its content is very dangerous.                                 |
| 5     | Extreme    | The message’s content is confirmed to part of an outbreak that is either extremely large scale or large scale and extremely dangerous. |

For more information about threat levels and outbreak rules, see [Outbreak Filters Rules, page 14-15](#).

**Related Topics**

- [Guidelines for Setting Your Quarantine Threat Level Threshold, page 14-8](#)
- [Containers: Specific and Always Rules, page 14-8](#)

## Guidelines for Setting Your Quarantine Threat Level Threshold

The quarantine threat level threshold allows administrators to be more or less aggressive in quarantining suspicious messages. A low setting (1 or 2) is more aggressive and will quarantine more messages; conversely, a higher score (4 or 5) is less aggressive and will only quarantine messages with an extremely high likelihood of being malicious.

The same threshold applies to both virus outbreaks and non-virus threats, but you can specify different quarantine retention times for virus attacks and other threats. See [Dynamic Quarantine, page 14-10](#) for more information.

Cisco recommends the default value of 3.

## Containers: Specific and Always Rules

Container files are files, such as zipped (.zip) archives, that contain other files. The TOC can publish rules that deal with specific files within archive files.

For example, if a virus outbreak is identified by TOC to consist of a .zip file containing a .exe, a specific Outbreak Rule is published that sets a threat level for .exe files within .zip files (.zip(exe)), but does not set a specific threat level for any other file type contained within .zip files (e.g. .txt files). A second rule (.zip(\*)) covers all other file types within that container file type. An Always rule for a container will always be used in a message's Threat Level calculation regardless of the types of files that are inside a container. An always rule will be published by the SIO if all such container types are known to be dangerous.

**Table 14-2**      *Fallback Rules and Threat Level Scores*

| Outbreak Rule | Threat Level | Description                                                                                           |
|---------------|--------------|-------------------------------------------------------------------------------------------------------|
| .zip(exe)     | 4            | This rule sets a threat level of 4 for .exe files within .zip files.                                  |
| .zip(doc)     | 0            | This rule sets a threat level of 0 for .doc files within .zip files.                                  |
| zip(*)        | 2            | This rule sets a threat level of 2 for all .zip files, regardless of the types of files they contain. |

## How the Outbreak Filters Feature Works

Email messages pass through a series of steps, the “email pipeline,” when being processed by your appliance (for more information about the email pipeline, see [Understanding the Email Pipeline, page 4-1](#)). As the messages proceed through the email pipeline, they are run through the anti-spam and anti-virus scanning engines if those engines are enabled for that mail policy. In other words, known spam or messages containing recognized viruses are not scanned by the Outbreak Filters feature because they will have already been removed from the mail stream — deleted, quarantined, etc. — based on your anti-spam and anti-virus settings. Messages that arrive at the Outbreak Filters feature have therefore been marked spam- and virus-free. Note that a message quarantined by Outbreak Filters may be marked as spam or containing a virus when it is released from the quarantine and rescanned by CASE, based on updated spam rules and virus definitions.

**Note**

Messages that skip anti-spam and anti-virus scanning due to filters or the engines being disabled will still be scanned by Outbreak Filters.

**Related Topics**

- [Message Scoring, page 14-9](#)
- [Dynamic Quarantine, page 14-10](#)

## Message Scoring

When a new virus attack or non-viral threat is released into the wild, no anti-virus or anti-spam software is able to recognize the threat yet, so this is where the Outbreak Filters feature can be invaluable. Incoming messages are scanned and scored by CASE using the published Outbreak and Adaptive Rules (see [Types of Rules: Adaptive and Outbreak, page 14-6](#)). The message score corresponds with the message's threat level. Based on which, if any, rules the message matches, CASE assigns the corresponding threat level. If there is no associated threat level (the message does not match any rules), then the message is assigned a threat level of 0.

Once that calculation has been completed, the Email Security appliance checks whether the threat level of that message meets or exceeds your quarantine or message modification threshold value and quarantines message or rewrites its URLs. If the threat level is below the thresholds, it will be passed along for further processing in the pipeline.

Additionally, CASE reevaluates existing quarantined messages against the latest rules to determine the latest threat level of a message. This ensures that only messages that have a threat level consistent with an outbreak message stay within the quarantine and messages that are no longer a threat flow out of the quarantine after an automatic reevaluation.

In the case of multiple scores for an outbreak message — one score from an Adaptive Rule (or the highest score if multiple Adaptive Rules apply), and another score from an Outbreak Rule (or the highest score if multiple Outbreak Rules apply) — intelligent algorithms are used to determine the final threat level.

**Note**

It is possible to use the Outbreak Filters feature without having enabled anti-virus scanning on the appliance. The two security services are designed to complement each other, but will also work separately. That said, if you do not enable anti-virus scanning on your appliance, you will need to monitor your anti-virus vendor's updates and manually release or re-evaluate some messages in the Outbreak quarantine. When using Outbreak Filters without anti-virus scanning enabled, keep the following in mind:

- You should disable Adaptive Rules
- Messages will get quarantined by Outbreak Rules
- Messages will get released if the threat level is lowered or time expires

Downstream anti-virus vendors (desktops/groupware) may catch the message on release.

**Note**

Anti-spam scanning needs to be enabled globally on an appliance in order for the Outbreak Filters feature to scan for non-viral threats.



## Dynamic Quarantine

The Outbreak Filters feature's Outbreak quarantine is a temporary holding area used to store messages until they're confirmed to be threats or it's safe to deliver to users. (See [Outbreak Lifecycle and Rules Publishing, page 14-11](#) for more information.) Quarantined messages can be released from the Outbreak quarantine in several ways. As new rules are downloaded, messages in the Outbreak quarantine are reevaluated based on a recommended rescan interval calculated by CASE. If the revised threat level of a message falls under the quarantine retention threshold, the message will automatically be released (regardless of the Outbreak quarantine's settings), thereby minimizing the time it spends in the quarantine. If new rules are published while messages are being re-evaluated, the rescan is restarted.

Please note that messages quarantined as virus attacks are not automatically released from the outbreak quarantine when new anti-virus signatures are available. New rules may or may not reference new anti-virus signatures; however, messages will not be released due to an anti-virus engine update unless an Outbreak Rule changes the threat level of the message to a score lower than your Threat Level Threshold.

Messages are also released from the Outbreak quarantine after CASE's recommended retention period has elapsed. CASE calculates the retention period based on the message's threat level. You can define separate maximum retention times for virus outbreaks and non-viral threats. If CASE's recommended retention time exceeds the maximum retention time for the threat type, the Email Security appliance releases messages when the maximum retention time elapses. For viral messages the default maximum quarantine period is 1 day. The default period for quarantining non-viral threats is 4 hours. You can manually release messages from the quarantine.

The Email Security appliance also releases messages when the quarantine is full and more messages are inserted (this is referred to as overflow). Overflow only occurs when the Outbreak quarantine is at 100% capacity, and a new message is added to the quarantine. At this point, messages are released in the following order of priority:

- Messages quarantined by Adaptive Rules (those scheduled to be released soonest are first)
- Messages quarantined by Outbreak Rules (those scheduled to be released soonest are first)

Overflow releases stop the moment the Outbreak quarantine is below 100% capacity. For more information about how quarantine overflow is handled, see [Retention Time for Messages in Quarantines, page 30-3](#) and [Default Actions for Automatically Processed Quarantined Messages, page 30-4](#).

Messages released from the Outbreak quarantine are scanned by the anti-virus and anti-spam engines again if they're enabled for the mail policy. If it is now marked as a known virus or spam, then it will be subject to your mail policy settings (including a possible second quarantining in the Virus quarantine or Spam quarantine). For more information, see [The Outbreak Filters Feature and the Outbreak Quarantine, page 14-21](#).

Thus it is important to note that in a message's lifetime, it may actually be quarantined twice — once due to the Outbreak Filters feature, and once when it is released from the Outbreak quarantine. A message will not be subject to a second quarantine if the verdicts from each scan (prior to Outbreak Filters, and when released from the Outbreak quarantine) match. Also note that the Outbreak Filters feature does not take any final actions on messages. The Outbreak Filters feature will either quarantine a message (for further processing) or move the message along to the next step in the pipeline.

### Related Topics

- [Outbreak Lifecycle and Rules Publishing, page 14-11](#)

## Outbreak Lifecycle and Rules Publishing

Very early in a virus outbreak's life cycle, broader rules are used to quarantine messages. As more information becomes available, increasingly focused rules are published, narrowing the definition of what is quarantined. As the new rules are published, messages that are no longer considered possible virus messages are released from quarantine (messages in the outbreak quarantine are rescanned as new rules are published).

**Table 14-3**      *Example Rules for an Outbreak Lifecycle*

| Time       | Rule Type                                  | Rule Description                                                                                                     | Action                                                                                       |
|------------|--------------------------------------------|----------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| T=0        | Adaptive Rule<br>(based on past outbreaks) | A consolidated rule set based on over 100K message attributes, which analyzes message content, context and structure | Messages are automatically quarantined if they match Adaptive Rules                          |
| T=5 min    | Outbreak Rule                              | Quarantine messages containing .zip (exe) files                                                                      | Quarantine all attachments that are .zips containing a .exe                                  |
| T=10 min   | Outbreak Rule                              | Quarantine messages that have .zip (exe) files greater than 50 KB                                                    | Any message with .zip (exe) files that are less than 50 KB would be released from quarantine |
| T=20 min   | Outbreak Rule                              | Quarantine messages that have .zip (exe) files between 50 to 55 KB, and have "Price" in the file name                | Any message that does not match this criteria would be released from quarantine              |
| T=12 hours | Outbreak Rule                              | Scan against new signature                                                                                           | All remaining messages are scanned against the latest anti-virus signature                   |

## Managing Outbreak Filters

Log in to the Graphical User Interface (GUI), select Security Services in the menu, and click Outbreak Filters.

**Figure 14-2 Outbreak Filters Main Page**  
**Outbreak Filters**

Outbreak Filters Overview

|                               |         |
|-------------------------------|---------|
| Global Status:                | Enabled |
| Adaptive Rules:               | Enabled |
| Maximum Message Size to Scan: | 512K    |
| Receive Emailed Alerts:       | No      |

Edit Global Settings...

Outbreak Filter Rules

| Rule Updates                                                                                      |                  |                                                                                                         |
|---------------------------------------------------------------------------------------------------|------------------|---------------------------------------------------------------------------------------------------------|
| Rule Type                                                                                         | Last Update      | Current Version                                                                                         |
| CASE Core Files                                                                                   | Never Updated    | 3.1.0-012                                                                                               |
| CASE Utilities                                                                                    | Never Updated    | 3.1.0-012                                                                                               |
| Virus Outbreak Rules                                                                              | Never Updated    | 20050718_000000                                                                                         |
| Outbreak Filter Rules (higher number indicates greater risk. 1= lowest threat, 5= highest threat) |                  |                                                                                                         |
| 3                                                                                                 | OUTBREAK_0003427 | We are seeing unusual volume for file extension(s) pif. We are raising the Threat Level to 3. We wil... |
| 3                                                                                                 | OUTBREAK_0003428 | We are seeing unusual volume for file extension(s) exe. We are raising the Threat Level to 3. We wil... |
| 3                                                                                                 | OUTBREAK_0003429 | We are seeing unusual volume for file extension(s) zip(exe), zip:e(exe). We are raising the Threat L... |
| 3                                                                                                 | OUTBREAK_0003430 | We are seeing suspicious url(s) propagating through multiple sources. We are raising the Threat Leve... |
| 3                                                                                                 | OUTBREAK_0003431 | We are seeing suspicious url(s) propagating through multiple sources. We are raising the Threat Leve... |

Rules last updated: Wed May 25 22:36:12 2011

Update Rules Now

Clear Current Rules

The Outbreak Filters page shows two sections: the Outbreak Filters Overview and a listing of current Outbreak Filter Rules (if any).

In [Figure 14-2](#), Outbreak Filters are enabled, Adaptive Scanning is enabled, and the maximum message size is set to 512k. To change these settings, click **Edit Global Settings**. For more information about editing Global Settings, see [Configuring Outbreak Filters Global Settings, page 14-12](#).

The Outbreak Filter Rules section lists the time, date, and version of the latest update for various components (the rules engine as well as the rules themselves), as well as a listing of the current Outbreak Filter rules with threat level.

For more information about Outbreak Rules, see [Outbreak Filters Rules, page 14-15](#).

#### Related Topics

- [Configuring Outbreak Filters Global Settings, page 14-12](#)
- [Outbreak Filters Rules, page 14-15](#)
- [The Outbreak Filters Feature and Mail Policies, page 14-15](#)
- [The Outbreak Filters Feature and the Outbreak Quarantine, page 14-21](#)

## Configuring Outbreak Filters Global Settings

To configure the Global Settings for Outbreak Filters, click **Edit Global Settings**.

**Figure 14-3** Outbreak Filters Global Settings Page  
Edit Outbreak Filters Settings

| Outbreak Filters Global Settings                            |                                                                                                      |
|-------------------------------------------------------------|------------------------------------------------------------------------------------------------------|
| <input checked="" type="checkbox"/> Enable Outbreak Filters |                                                                                                      |
| Adaptive Rules:                                             | <input checked="" type="checkbox"/> Enable Adaptive Rules                                            |
| Maximum Message Size to Scan:                               | <input type="text" value="512k"/> Maximum<br><small>Add a trailing K or M to indicate units.</small> |
| Emailed Alerts: ?                                           | <input type="checkbox"/> Receive Emailed Alerts                                                      |

Use this page to:

- Enable Outbreak Filters globally
- Enable Adaptive Rules scanning
- Set a maximum size for files to scan (note that you are entering the size in *bytes*)
- Enable alerts for the Outbreak Filter

Note that alerts and Adaptive Rules are not enabled by default. This functionality is also available via the `outbreakconfig` CLI command (see the *Cisco AsyncOS CLI Reference Guide*). After you make your changes, submit and commit them.



**Note**

You cannot enable the logging of URLs using the web interface. For instructions to enable logging of URLs using CLI, see [Enabling Logging of URLs, page 14-14](#).

## Enabling the Outbreak Filters Feature

To enable the Outbreak Filters feature globally, check the box next to Enable Outbreak Filters on the Outbreak Filters Global Settings page, and click **Submit**. You must have agreed to the Outbreak Filters license agreement first.

Once enabled globally, the Outbreak Filters feature can then be enabled or disabled individually for each incoming and outgoing mail policy, including the default policies. For more information, see [The Outbreak Filters Feature and Mail Policies, page 14-15](#).

The Outbreak Filters feature uses the Context Adaptive Scanning Engine (CASE) to detect viral threats, regardless of whether anti-spam scanning is enabled, but you do need to have Anti-Spam or Intelligent Multi-Scan enabled globally on the appliance in order to scan for non-viral threats.



**Note**

If you have not already agreed to the license during system setup (see [Step 4: Security, page 3-21](#)), you must click **Enable** on the Security Services > Outbreak Filters page, and then read and agree to the license.

## Enabling Adaptive Rules

Adaptive Scanning enables the use of Adaptive Rules in Outbreak Filters. A set of factors or traits (file size, etc.) are used to determine the likelihood of a message being part of an outbreak when no virus signature or spam criteria relating to the message's content is available. To enable Adaptive Scanning, check the box next to Enable Adaptive Rules on the Outbreak Filters Global Settings page, and click **Submit**.

## Enabling Alerts for Outbreak Filters

Check the box labeled “Emailed Alerts” to enable alerting for the Outbreak Filters feature. Enabling emailed alerts for Outbreak Filters merely enables the alerting engine to send alerts regarding Outbreak Filters. Specifying which alerts are sent and to which email addresses is configured via the Alerts page in the System Administration tab. For more information on configuring alerts for Outbreak Filters, see [Alerts, SNMP Traps, and Outbreak Filters, page 14-23](#).

## Enabling Logging of URLs

Logging of URL-related logs is disabled by default. This includes the logs for the following events:

- Category of any URL in the message matches the URL category filters
- Reputation score of any URL in the message matches URL reputation filters
- Outbreak Filter rewrites any URL in the message

Use the `outbreakconfig` command in CLI to enable the logging of these events.

### Related Topics

- [Example, page 14-14](#)

## Example

The following example shows how to enable logging of URLs using the `outbreakconfig` command

```
mail.example.com> outbreakconfig
```

```
Outbreak Filters: Enabled
```

```
Choose the operation you want to perform:
```

```
- SETUP - Change Outbreak Filters settings.
```

```
[> setup
```

```
Outbreak Filters: Enabled
```

```
Would you like to use Outbreak Filters? [Y]>
```

```
Outbreak Filters enabled.
```

```
Outbreak Filter alerts are sent when outbreak rules cross the threshold (go above or back down below), meaning that new messages of certain types could be quarantined or will no longer be quarantined, respectively.
```

```
Would you like to receive Outbreak Filter alerts? [N]>
```

```
What is the largest size message Outbreak Filters should scan?
```

```
[524288]>
```

```
Do you want to use adaptive rules to compute the threat level of messages? [Y]>
```

```
Logging of URLs is currently disabled.
```

```
Do you wish to enable logging of URL's? [N]> Y
```

```
Logging of URLs has been enabled.
```

```
The Outbreak Filters feature is now globally enabled on the system. You must use the 'policyconfig' command in the CLI or the Email Security Manager in the GUI to enable Outbreak Filters for the desired Incoming and Outgoing Mail Policies.
```

```
Choose the operation you want to perform:
- SETUP - Change Outbreak Filters settings.
[]>
```

## Outbreak Filters Rules

Outbreak Rules are published by the Cisco Security Intelligence Operations and your appliance checks for and downloads new outbreak rules every 5 minutes. You can change this update interval. See [Configuring Server Settings for Downloading Upgrades and Updates, page 33-21](#) for more information.

### Related Topics

- [Managing Outbreak Filter Rules, page 14-15](#)

## Managing Outbreak Filter Rules

Because the Outbreak Filters Rules are automatically downloaded for you, there really is no management needed on the part of the user.

However, if for some reason your appliance is not able to reach Cisco's update servers for new rules over a period of time, it is possible that your locally-cached scores are no longer valid, i.e., if a known viral attachment type now has an update in the anti-virus software and/or is no longer a threat. At this time, you may wish to no longer quarantine messages with these characteristics.

You can manually download updated outbreak rules from Cisco's update servers by clicking **Update Rules Now**.



### Note

The **Update Rules Now** button does not “flush” all existing outbreak rules on the appliance. It only replaces outbreak rules that have been updated. If there are no updates available on Cisco's update servers, then the appliance will not download any outbreak rules when you click this button.

### Related Topics

- [Updating Outbreak Filter Rules, page 14-15](#)

## Updating Outbreak Filter Rules

By default, your appliance will attempt to download new Outbreak Filters rules every 5 minutes. You can change this interval via the Security Services > Service Updates page. For more information, see [Service Updates, page 33-17](#).

## The Outbreak Filters Feature and Mail Policies

The Outbreak Filters feature has settings that can be set per mail policy. The Outbreak Filters feature can be enabled or disabled for each mail policy on the appliance. Specific file extensions and domains can be exempted from processing by the Outbreak Filters feature, per mail policy. This functionality is also available via the `policyconfig` CLI command (see the *Cisco AsyncOS CLI Reference Guide*).



### Note

Anti-Spam or Intelligent Multi-Scan scanning needs to be enabled globally on an appliance in order for the Outbreak Filters feature to scan for non-viral threats.

**Figure 14-4 Mail Policy Listing**  
Incoming Mail Policies

**Find Policies**

Email Address:

☒ Recipient ☐ Sender [Find Policies](#)

**Policies**

[Add Policy...](#)

| Order | Policy Name    | Anti-Spam                                                                                       | Anti-Virus                                                                   | Content Filters                                 | Outbreak Filters                                  | Delete |
|-------|----------------|-------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|-------------------------------------------------|---------------------------------------------------|--------|
| 1     | Sales_Team     | IronPort Anti-Spam<br>Positive: Drop<br>Suspected: Quarantine<br>Marketing Messages: Quarantine | (use default)                                                                | (use default)                                   | (use default)                                     |        |
| 2     | Engineering    | (use default)                                                                                   | (use default)                                                                | scan_for_confidential<br>ex_employee            | Retention Time:<br>Virus: 1 day<br>Other: 4 hours |        |
|       | Default Policy | IronPort Anti-Spam<br>Positive: Drop<br>Suspected: Deliver<br>Marketing Messages: Disabled      | Sophos<br>Encrypted: Deliver<br>Unscannable: Deliver<br>Virus Positive: Drop | scan_for_confidential<br>no_mp3s<br>ex_employee | Retention Time:<br>Virus: 1 day                   |        |

Key: Default Custom Readonly

To modify the Outbreak Filters feature settings for a specific mail policy, click the link in the Outbreak Filters column of the policy to change.

**Figure 14-5 Outbreak Filters Settings and Mail Policies**

**Outbreak Filtering for: Default Policy**

Enable Outbreak Filtering (Customize settings) [▼](#)

**Outbreak Filter Settings**

Quarantine Threat Level: [?](#)

Maximum Quarantine Retention:

Viral Attachments:  Days [▼](#)

Other Threats:  Hours [▼](#)

☐ Deliver messages without adding them to quarantine

Bypass Attachment Scanning: [▶](#) None configured

**Message Modification**

☒ Enable message modification. Required for non-viral threat detection (excluding attachments)

Message Modification Threat Level: [?](#)

Message Subject: [Prepend](#)  [Insert Variables](#) | [Preview Text](#) [↗](#)

Include the X-IronPort-Outbreak headers:

☒ Enable for all messages

☐ Enable only for threat-based outbreak

☐ Disable

Include the X-IronPort-Outbreak-Description header:

☒ Enable

☐ Disable

Alternate Destination Mail Host:

(examples: example.com, 10.0.0.1, 2001:420:80:1::5)

URL Rewriting: Cisco Security proxy scans and rewrites all URLs contained in malicious outbreak emails.

☒ Enable only for unsigned messages (recommended)

☐ Enable for all messages

☐ Disable

Bypass Domain Scanning [?](#)

(examples: example.com, crm.example.com, 10.0.0.1, 10.0.0.0/24, 2001:420:80:1::5, 2001:db8::/32)

Threat Disclaimer: [None](#) [▼](#)

Disclaimer text will be applied to the top of the message body for Suspicious and Quarantined messages. To create custom disclaimers go to Mail Policies > Text Resources > Disclaimers

To enable and customize the Outbreak Filters feature for a particular mail policy, select **Enable Outbreak Filtering (Customize Settings)**.

You can configure the following Outbreak Filter settings for a mail policy:

- Quarantine threat level
- Maximum quarantine retention time
- Deliver non-viral threat messages immediately without adding them to quarantine
- File extension types for bypassing
- Message modification threshold
- Alter subject header using custom text and Outbreak Filter variables such as `$threat_verdict`, `$threat_category`, `$threat_type`, `$threat_description`, and `$threat_level`.
- Include the following email headers:
  - X-IronPort-Outbreak-Status
  - X-IronPort-Outbreak-Description
- Send the message to an alternate destination such as an Email Security Appliance or an exchange server.
- URL rewriting
- Threat disclaimer

Select **Enable Outbreak Filtering (Inherit Default mail policy settings)** to use the Outbreak Filters settings that are defined for the default mail policy. If the default mail policy has the Outbreak Filters feature enabled, all other mail policies use the same Outbreak Filter settings unless they are customized.

Once you have made your changes, commit your changes.

#### Related Topics

- [Setting a Quarantine Level Threshold, page 14-17](#)
- [Maximum Quarantine Retention, page 14-17](#)
- [Bypassing File Extension Types, page 14-18](#)
- [Message Modification, page 14-18](#)

## Setting a Quarantine Level Threshold

Select a Quarantine Threat Level threshold for outbreak threats from the list. A smaller number means that you will be quarantining more messages, while a larger number results in fewer messages quarantined. Cisco recommends the default value of 3.

For more information, see [Guidelines for Setting Your Quarantine Threat Level Threshold, page 14-8](#).

## Maximum Quarantine Retention

Specify the maximum amount of time in either hours or days that messages stay in the Outbreak Quarantine. You can specify different retention times for messages that may contain viral attachments and messages that may contain other threats, like phishing or malware links. For non-viral threats, check the **Deliver messages without adding them to quarantine** check box to deliver the messages immediately without adding them to quarantine.



**Note**

You cannot quarantine non-viral threats unless you enable Message Modification for the policy.

CASE recommends a quarantine retention period when assigning the threat level to the message. The Email Security appliance keeps the message quarantined for the length of time that CASE recommends unless it exceeds the maximum quarantine retention time for its threat type.

## Bypassing File Extension Types

You can modify a policy to bypass specific file types. Bypassed file extensions are not included when CASE calculates the threat level for the message; however, the attachments are still processed by the rest of the email security pipeline.

To bypass a file extension, click Bypass Attachment Scanning, select or type in a file extension, and click **Add Extension**. AsyncOS displays the extension type in the File Extensions to Bypass list.

To remove an extension from the list of bypassed extensions, click the trash can icon next to the extension in the File Extensions to Bypass list.

**Related Topics**

- [Bypassing File Extensions: Container File Types, page 14-18](#)

### Bypassing File Extensions: Container File Types

When bypassing file extensions, files within container files (a .doc file within a .zip, for example) are bypassed if the extension is in the list of extensions to bypass. For example, if you add .doc to the list of extensions to bypass, all .doc files, even those within container files are bypassed.

## Message Modification

Enable Message Modification if you want the appliance to scan messages for non-viral threats, such as phishing attempts or links to malware websites.

Based on the message's threat level, AsyncOS can modify the message to rewrite all of the URLs to redirect the recipient through the Cisco web security proxy if they attempt to open the website from the message. The appliance can also add a disclaimer to the message to alert the user that the message's content is suspicious or malicious.

You need to enable message modification in order to quarantine non-viral threat messages.

**Related Topics**

- [Message Modification Threat Level, page 14-19](#)
- [Message Subject, page 14-19](#)
- [Outbreak Filters Email Headers, page 14-19](#)
- [Alternate Destination Mail Host, page 14-19](#)
- [URL Rewriting and Bypassing Domains, page 14-20](#)
- [Threat Disclaimer, page 14-20](#)

## Message Modification Threat Level

Select a Message Modification Threat Level threshold from the list. This setting determines whether to modify a message based on the threat level returned by CASE. A smaller number means that you will be modifying more messages, while a larger number results in fewer messages being modified. Cisco recommends the default value of 3.

## Message Subject

You can alter the text of the subject header on non-viral threat messages containing modified links to notify users that the message has been modified for their protection. Prepend or append the subject header with custom text, Outbreak Filter variables such as `$threat_verdict`, `$threat_category`, `$threat_type`, `$threat_description`, and `$threat_level`, or a combination of both. To insert variables, click **Insert Variables**, and select from the list of variables.

White space is not ignored in the Message Subject field. Add spaces after (if prepending) or before (if appending) the text you enter in this field to separate your added text from the original subject of the message. For example, add the text `[MODIFIED FOR PROTECTION]` with a few trailing spaces if you are prepending.



### Note

The Message Subject field only accepts US-ASCII characters.

## Outbreak Filters Email Headers

You can add the following additional headers to the message:

| Header                                 | Format                                                                                                         | Example                                                                                                              | Options                                                                                                                                      |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>X-IronPort-Outbreak-Status</b>      | X-IronPort-Outbreak-Status:<br>\$threat_verdict, level<br>\$threat_level, \$threat_category<br>- \$threat_type | X-IronPort-Outbreak-Status: Yes, level 4, Phish<br>- Password                                                        | <ul style="list-style-type: none"> <li>• Enable for all messages</li> <li>• Enable only for non-viral outbreak</li> <li>• Disable</li> </ul> |
| <b>X-IronPort-Outbreak-Description</b> | X-IronPort-Outbreak-Description:<br>\$threat_description                                                       | X-IronPort-Outbreak-Description: It may trick victims into submitting their username and password on a fake website. | <ul style="list-style-type: none"> <li>• Enable</li> <li>• Disable</li> </ul>                                                                |



### Note

If you want to filter messages based on these headers, you must send the Outbreak Filter processed messages back to an Email Security Appliance (by configuring an alternate destination mail host), and scan them using a content filter that matches these headers.

## Alternate Destination Mail Host

If you want to perform a content filter-based scan on the Outbreak Filter processed messages, you must configure the Outbreak Filter to send the processed messages back to an Email Security Appliance. This is because, in the processing pipeline, the Outbreak Filter scan is performed after the content filter scan.

In the **Alternate Destination Mail Host** field, enter the IP address (IPv4 or IPv6) or the FQDN of the appliance where you want to send the processed messages for further scans.

## URL Rewriting and Bypassing Domains

If the message's threat level exceeds the message modification threshold, the Outbreak Filters feature rewrites all URLs in the message to redirect the user to the Cisco web security proxy's splash page if they click on any of them. (See [Redirecting URLs, page 14-5](#) for more information.) If the message's threat level exceeds the quarantine threshold, the appliance also quarantines the message. If a small scale, non-viral outbreak is in progress, quarantining the message gives TOC time to analyze any suspect websites linked from possible outbreak messages and determine whether the websites are malicious. CASE uses updated Outbreak Rules from SIO to rescan the message to determine if it is part of the outbreak. After the retention period expires, the appliance releases the message from the quarantine.

AsyncOS rewrites all of the URLs inside a message except for the ones pointing to bypassed domains.

The following options are available for URL rewriting:

- **Enable only for unsigned messages.** This option allows AsyncOS to rewrite URLs in unsigned messages that meet or exceed the message modification threshold, but not signed messages. Cisco recommends using this setting for URL rewriting.



### Note

The Email Security appliance may rewrite URLs in a DomainKeys/DKIM-signed message and invalidate the message's signature if a server or appliance on your network other than the Email Security appliance is responsible for verifying the DomainKeys/DKIM signature.

- **Enable for all messages.** This option allows AsyncOS to rewrite URLs in all messages that meet or exceed the message modification threshold, including signed ones. If AsyncOS modifies a signed message, the signature becomes invalid.
- **Disable.** This option disables URL rewriting for Outbreak Filters.

You can modify a policy to exclude URLs to certain domains from modification. To bypass domains, enter the IPv4 address, IPv6 address, CIDR range, hostname, partial hostname or domain in the Bypass Domain Scanning field. Separate multiple entries using commas.

The Bypass Domain Scanning feature is similar to, but independent of, the global whitelist used by URL filtering. For more information about that whitelist, see [Creating Whitelists for URL Filtering, page 15-4](#).

## Threat Disclaimer

The Email Security appliance can append a disclaimer message above the heading of a suspicious message to warn the user of its content. This disclaimer can be in HTML or plain text, depending on the type of message.

Select the disclaimer text you want to use from the Threat Disclaimer list or click the Mail Policies > Text Resources link to create a new disclaimer using the Disclaimer Template. The Disclaimer Template includes variables for outbreak threat information. You can see a preview of the threat disclaimer by clicking Preview Disclaimer. For custom disclaimer messages, you can use variables to display the threat level, the type of threat, and a description of the threat in the message. For information on creating a disclaimer message, see [Overview of Text Resource Management, page 21-9](#).

## The Outbreak Filters Feature and the Outbreak Quarantine

Messages quarantined by the Outbreak Filters feature are sent to the Outbreak quarantine. This quarantine functions like any other quarantine (for more information about working with quarantines, see [Chapter 30, “Policy, Virus, and Outbreak Quarantines”](#)) except that it has a “summary” view, useful for deleting or releasing all messages from the quarantine, based on the rule used to place the message in the quarantine (for Outbreak Rules, the Outbreak ID is shown, and for Adaptive Rules, a generic term is shown). For more information about the summary view, see [Outbreak Quarantine and the Manage by Rule Summary View](#), page 14-22.

### Related Topics

- [Monitoring the Outbreak Quarantine](#), page 14-21
- [Outbreak Quarantine and the Manage by Rule Summary View](#), page 14-22

## Monitoring the Outbreak Quarantine

Though a properly configured quarantine requires little if any monitoring, it is a good idea to keep an eye on the Outbreak Quarantine, especially during and after virus outbreaks when legitimate messages may be delayed.

If a legitimate message is quarantined, one of the following occurs depending on the settings for the Outbreak quarantine:

- If the quarantine’s Default Action is set to Release, the message will be released when the retention time period expires or when the quarantine overflows. You can configure the Outbreak quarantine so that the following actions are performed on messages before they are released due to overflow: strip attachments, modify the subject, and add an X-Header. For more information about these actions, see [Default Actions for Automatically Processed Quarantined Messages](#), page 30-4.
- If the quarantine’s Default Action is set to Delete, the message will be deleted when the retention time period expires, or when the quarantine overflows.
- Overflow occurs when the quarantine is full and more messages are added. In this case the messages closest to their expiration date (not necessarily the oldest messages) are released first, until enough room is available for the new messages. You can configure the Outbreak quarantine so that the following actions are performed on messages before they are released due to overflow: strip attachments, modify the subject, add an X-Header.

Because quarantined messages are rescanned whenever new rules are published, it is very likely that messages in the Outbreak quarantine will be released prior to the expiration time.

Still, it can be important to monitor the Outbreak quarantine if the Default Action is set to Delete. Cisco recommends most users to not set the default action to Delete. For more information about releasing messages from the Outbreak quarantine, or changing the Default Action for the Outbreak Quarantine, see [Default Actions for Automatically Processed Quarantined Messages](#), page 30-4.

Conversely, if you have messages in your Outbreak quarantine that you would like to keep in the quarantine longer while you wait for a new rule update, for example, you can delay the expiration of those messages. Keep in mind that increasing the retention time for messages can cause the size of the quarantine to grow.



### Note

If anti-virus scanning is disabled globally (not via a mail policy) while a message is in the Outbreak quarantine, the message is not anti-virus scanned when it leaves the quarantine, even if anti-virus scanning is re-enabled prior to the message leaving the quarantine.

**Note**

You can use the Outbreak Filters feature without having enabled anti-virus scanning on the appliance. However, Outbreak Filters cannot scan for non-viral threats if anti-spam scanning is not enabled on the appliance.

## Outbreak Quarantine and the Manage by Rule Summary View

You can view the contents of the Outbreak quarantine by clicking on the name of the quarantine in the listing on the Monitor menu in the GUI. The Outbreak quarantine has an additional view as well, the Outbreak Quarantine Manage by Rule Summary link.

**Figure 14-6** The Outbreak Quarantine Manage by Rule Summary Link  
**Quarantines**

| Quarantines                                          |          |                                     |                                |                      |
|------------------------------------------------------|----------|-------------------------------------|--------------------------------|----------------------|
| <a href="#">Add Quarantine...</a>                    |          |                                     |                                |                      |
| Quarantine                                           | Messages | Default Action                      | Status                         | Settings             |
| Spam Quarantine                                      | 2565     | Retain 14 days then Delete          | <div><div></div></div> 2% Full | <a href="#">Edit</a> |
| Outbreak<br><a href="#">[Manage by Rule Summary]</a> | 0        | Retention Varies<br>Action: Release | <div><div></div></div> 0% Full | <a href="#">Edit</a> |
| Policy                                               | 0        | Retain 10 days then Delete          | <div><div></div></div> 0% Full | <a href="#">Edit</a> |
| Virus                                                | 0        | Retain 30 days then Delete          | <div><div></div></div> 0% Full | <a href="#">Edit</a> |

### Related Topics

- [Using the Summary View to Perform Message Actions on Messages in the Outbreak Quarantine Based on Rule ID.](#), page 14-22

### Using the Summary View to Perform Message Actions on Messages in the Outbreak Quarantine Based on Rule ID.

Click on the Manage by Rule Summary link to see a listing of the contents of the Outbreak quarantine, grouped by rule ID:

**Figure 14-7** The Outbreak Quarantine Manage by Rule Summary View  
**Outbreak Quarantine Summary**

| Manage by Rule Summary                                                                |          |                    |                      |            |          |
|---------------------------------------------------------------------------------------|----------|--------------------|----------------------|------------|----------|
| <input type="checkbox"/> All<br>Select                                                | Rule ID  | Number of messages | Average message size | Total size | Capacity |
| <input type="checkbox"/>                                                              | EXE_BAGL | 4                  | 16 KB                | 0.1 MB     | 0.0%     |
| <b>Totals</b>                                                                         |          | 4                  | 16 KB                |            |          |
| <input type="button" value="Select Action..."/> <input type="button" value="Submit"/> |          |                    |                      |            |          |

From this view, you can choose to release, delete, or delay the exit for all messages pertaining to a specific outbreak or adaptive rule, rather than selecting individual messages. You can also search through or sort the listing.

This functionality is also available via the `quarantineconfig -> outbreakmanage` CLI command. For more information, see the *Cisco AsyncOS CLI Reference Guide*.

# Monitoring Outbreak Filters

The appliance includes several tools to monitor the performance and activity of the Outbreak Filters feature.

## Related Topics

- [Outbreak Filters Report, page 14-23](#)
- [Outbreak Filters Overview and Rules Listing, page 14-23](#)
- [Outbreak Quarantine, page 14-23](#)
- [Alerts, SNMP Traps, and Outbreak Filters, page 14-23](#)

## Outbreak Filters Report

The Outbreak Filters report to view the current status and configuration of Outbreak Filters on your appliance as well as information about recent outbreaks and messages quarantined due to Outbreak Filters. View this information on the Monitor > Outbreak Filters page. For more information, see the “Email Security Monitor” chapter.

## Outbreak Filters Overview and Rules Listing

The overview and rules listing provide useful information about the current status of the Outbreak Filters feature. View this information via the Security Services > Outbreak Filters page.

## Outbreak Quarantine

Use the outbreak quarantine to monitor how many messages are being flagged by your Outbreak Filters threat level threshold. Also available is a listing of quarantined messages by rule. For information, see [Outbreak Quarantine and the Manage by Rule Summary View, page 14-22](#) and [Chapter 30, “Policy, Virus, and Outbreak Quarantines.”](#)

## Alerts, SNMP Traps, and Outbreak Filters

The Outbreak Filters feature supports two different types of notifications: regular AsyncOS alerts and SNMP traps.

SNMP traps are generated when a rule update fails. For more information about SNMP traps in AsyncOS, see the “Managing and Monitoring via the CLI” chapter.

AsyncOS has two types of alerts for the Outbreak Filter feature: size and rule

AsyncOS alerts are generated whenever the Outbreak quarantine’s size goes above 5, 50, 75, and 95 of the maximum size. The alert generated for the 95% threshold has a severity of CRITICAL, while the remaining alert thresholds are WARNING. Alerts are generated when the threshold is crossed as the quarantine size increases. Alerts are not generated when thresholds are crossed as the quarantine size decreases. For more information about alerts, see [Alerts, page 33-34](#).

AsyncOS also generates alerts when rules are published, the threshold changes, or when a problem occurs while updating rules or the CASE engine.

## Troubleshooting The Outbreak Filters Feature

This section provides some basic troubleshooting tips for the Outbreak Filters feature.

### Related Topics

- [Reporting Incorrectly Classified Messages to Cisco](#), page 14-24
- [Multiple Attachments and Bypassed Filetypes](#), page 14-24
- [Message and Content Filters and the Email Pipeline](#), page 14-24

## Reporting Incorrectly Classified Messages to Cisco

Use the checkbox on the Manage Quarantine page for the Outbreak quarantine to notify Cisco of misclassifications.

## Multiple Attachments and Bypassed Filetypes

Bypassed file types are only excluded if a message's only attachment is of that type, or in the case of multiple attachments, if the other attachments do not yet have existing rules. Otherwise the message is scanned.

## Message and Content Filters and the Email Pipeline

Message and content filters are applied to messages prior to scanning by Outbreak Filters. Filters can cause messages to skip or bypass the Outbreak Filters scanning.



## URL Filtering

---

- [Overview of URL Filtering, page 15-1](#)
- [Setting Up URL Filtering, page 15-2](#)
- [Taking Action Based on the Reputation or Category of URLs in Messages, page 15-7](#)
- [Monitoring URL Filtering Results, page 15-10](#)
- [Troubleshooting URL Filtering, page 15-10](#)
- [About URL Categories, page 15-13](#)

### Overview of URL Filtering

URL filtering uses the reputation and category of URL links in messages to:

- Increase the effectiveness of protection from malicious URLs in messages

URL filtering is incorporated into Outbreak Filtering. This strengthened protection is useful even if your organization already has a Cisco Web Security Appliance or similar protection from web-based threats, because it blocks threats at the point of entry.

You can also use content or message filters to take action based on the Web Based Reputation Score (WBRS) of URLs in messages. For example, you can rewrite URLs with suspect or unknown reputation to redirect them to the Cisco Web Security Proxy for click-time evaluation of their safety.

- Better identify spam

The appliance uses the reputation and category of links in messages, in conjunction with other spam-identification algorithms, to help identify spam. For example, if a link in a message belongs to a marketing web site, the message is more likely to be a marketing message.

- Support enforcement of corporate acceptable use policies

The category of URLs (for example, Adult Content or Illegal Activities) can be used in conjunction with content and message filters to enforce corporate acceptable use policies.

URL filtering is incorporated into the anti-spam, outbreak, content, and message filtering processes in the work queue.

### Which URLs Are Evaluated

URLs in incoming and outgoing messages (excluding attachments) are evaluated. Any valid string for a URL is evaluated, including strings with the following:



- http, https, or www
- domain or IP address
- port number preceded by a colon (:)
- uppercase or lowercase letters

When evaluating URLs to determine whether a message is spam, if necessary for load management, the system prioritizes screening of incoming messages over outgoing messages.

## Setting Up URL Filtering

### Related Topics

- [Requirements for URL Filtering, page 15-2](#)
- [Enable URL Filtering, page 15-2](#)
- [About the Connection to Cisco Web Security Services, page 15-3](#)
- [URL Filtering in Cluster Configurations, page 15-4](#)
- [Creating Whitelists for URL Filtering, page 15-4](#)

## Requirements for URL Filtering

In addition to enabling URL filtering, you must enable other features depending on desired functionality.

For enhanced protection against spam:

- Anti-spam scanning must be enabled globally and per applicable mail policy. This can be either the IronPort Anti-Spam or the Intelligent Multi-Scan feature. See the anti-spam chapter.

For enhanced protection against malware:

- The Outbreak Filters feature must be enabled globally and per applicable mail policy. See the Outbreak Filters chapter.

To take action based on URL reputation, or to enforce acceptable use policies using message and content filters:

- The Outbreak Filters feature must be enabled globally. See the Outbreak Filters chapter.

## Enable URL Filtering

### Before You Begin

- Ensure that the requirements for the individual URL filtering features that you want to use have been met. See [Requirements for URL Filtering, page 15-2](#).
- (Optional) Create a list of URLs that you want all URL filtering functionalities to ignore. See [Creating Whitelists for URL Filtering, page 15-4](#).

### Procedure

---

**Step 1** Select **Security Services > URL Filtering**.

- Step 2** Click **Enable**.
- Step 3** Select the **Enable URL Category and Reputation Filters** check box.
- Step 4** (Optional) If you have created a list of URLs to exempt from URL filtering when evaluating messages for spam and malware, and from all content and message filtering, select that list.
- This setting does not cause the message to bypass anti-spam or Outbreak Filters processing generally.
- Step 5** Submit and commit your changes.
- If you have met the applicable prerequisites, and you have already configured Outbreak Filters and Anti-Spam protection, then you do not need to make additional configurations to benefit from enhanced automatic detection of spam and malicious URLs.

#### What To Do Next

- To take action based on the reputation of URLs in messages, see [Taking Action Based on the Reputation or Category of URLs in Messages, page 15-7](#).
- To use URL categories in content and message filters, for example to enforce acceptable use policies, see [Taking Action Based on the Reputation or Category of URLs in Messages, page 15-7](#).
- To redirect all URLs in suspected spam messages to the Cisco Web Security proxy service, see [Using Custom Headers to Redirect URLs in Suspected Spam to the Cisco Web Security Proxy: Configuration Example, page 13-11](#).
- (Optional) To customize the appearance of end user notification page, see [Customizing the Appearance of End User Notification Page, page 15-6](#).
- Ensure that you receive alerts about issues related to this feature. See [Future URL Category Set Changes, page 15-22](#), the release notes for your AsyncOS release, and [Adding Alert Recipients, page 33-36](#).

## About the Connection to Cisco Web Security Services

URL reputation and category are provided by cloud-based Cisco Web Security Services.

The Email Security appliance connects to the Cisco Web Security Services either directly or through a web proxy, using the port specified for URL filtering services in [Appendix D, “Firewall Information.”](#) Communication is over HTTPS with mutual certificate authentication. Certificates are updated automatically (see [Service Updates, page 33-17](#).) For additional information about required certificates, see the Release Notes available from the location specified in [Certificates for URL Filtering Features, page 15-4](#).

If an HTTP or HTTPS proxy has been configured on the Security Services > Service Updates page, the Email Security appliance will use it when communicating with Cisco Web Security Services. For more information about using a proxy server, see [Configuring Server Settings for Downloading Upgrades and Updates, page 33-21](#).

In FIPS mode, communications with the Cisco Web Security Services uses FIPS ciphers.



#### Note

Certificates are not saved with a configuration file.

#### Related Topics

- [Certificates for URL Filtering Features, page 15-4](#)

- [Alert: SDS: Error Fetching Enrollment Certificate, page 15-11](#)
- [Alert: SDS: Certificate Is Invalid, page 15-11](#)

## Certificates for URL Filtering Features

AsyncOS is designed to automatically deploy and update the certificates needed for communications with cloud services used for URL filtering features. However, if for any reason the system is unable to update these certificates, you will receive an alert that requires action from you.

Ensure that the appliance is configured to send you these alerts (System type, Warning severity). For instructions, see [Alerts, page 33-34](#).

If you receive an alert about an invalid certificate, contact Cisco TAC, which can provide the required replacement certificate. For instructions to use the replacement certificate, see [Manually Configuring a Certificate for Communication with Cisco Web Security Services, page 15-13](#).

## URL Filtering in Cluster Configurations

- You can enable URL filtering at the machine, group or cluster level.
- If URL filtering is enabled at machine level, URL whitelists can be configured at machine, group or cluster level.
- If URL filtering is enabled at group level, URL whitelists must be configured at group or cluster level.
- If URL filtering is enabled at cluster level, URL whitelists must be configured at cluster level.
- The standard rules for clusters for Message Filters and Content Filters apply.

## Creating Whitelists for URL Filtering

If you specify a global whitelist when configuring the URL Filtering feature, then URLs on the whitelist are not evaluated for reputation or category, for anti-spam, Outbreak Filtering, or content and message filtering. However, the messages that contain these URLs are evaluated as usual by anti-spam scanning and Outbreak Filters. You can also specify a URL whitelist in each URL Filtering condition (rule) and action in content and message filters, to supplement the global URL whitelist.

To whitelist URLs from Outbreak Filtering generally, use the Bypass Domain Scanning option that you configure on the Mail Policies: Outbreak Filters page. URL whitelists for URL filtering are similar to, but independent of, Bypass Domain Scanning. For more information about that feature, see [URL Rewriting and Bypassing Domains, page 14-20](#).

There is no relationship between URL filtering whitelists described in this section and the whitelist used for sender reputation filtering based on SBRS score.

### Before You Begin

Consider importing a list of URLs instead of creating one in the web interface. See [Importing a URL List, page 15-5](#).

### Procedure

---

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

- Step 2** Select **Add URL List** or click a list to edit.
- Be sure all URLs that you want to globally whitelist are in a single list. You can select only one global whitelist for URL filtering.
- Step 3** Create and submit the URL list.
- To view a list of supported URL formats, enter a semicolon (;) into the **URLs** box and click **Submit**. Then click the **more...** link that appears.
- Each URL, domain, or IP address can be on a separate line, or separate each with a comma.
- Step 4** Commit your changes.
- 

**What To Do Next**

- To designate a URL list as the global whitelist, see [Enable URL Filtering, page 15-2](#).
- To designate a URL list as the whitelist for a specific condition (rule) or action in a content or message filter, see [Taking Action Based on the Reputation or Category of URLs in Messages, page 15-7](#) and [Content Filter Actions, page 11-9](#). For message filters, see also [URL Category Rule, page 9-47](#) and [URL Category Actions, page 9-76](#).

**Related Topics**

- [Importing a URL List, page 15-5](#)

## Importing a URL List

You can import a URL list to use as a whitelist for URL filtering.

**Procedure**

- 
- Step 1** Create the text file to import:
- The first line must be the name of the URL list.
  - Each URL must be on a separate line.
- Step 2** Upload the file to the `/configuration` directory on the appliance.
- Step 3** Use the `urllistconfig > new` command in the command-line interface.
- 

## Cisco Web Security Proxy End User Notification Page

If an end user clicks on a malicious URL rewritten by an Outbreak Filter or Policy (using Content or Message Filters), the Cisco Web Security proxy displays a Cisco branded notification page to the end user in the web browser. This page notifies that the site is malicious and access to it has been blocked. You can customize the appearance of this notification page and display your organization's branding such as company logo, contact information, and so on. See [Customizing the Appearance of End User Notification Page, page 15-6](#).

Keep in mind that, when an end user clicks on a URL rewritten using Outbreak Filter, the notification page is displayed for 10 seconds and then redirected to the Cisco Web Security proxy (Cisco branded) for click-time evaluation.

**Related Topics**

- [Redirecting URLs, page 14-5](#)
- [Taking Action Based on the Reputation or Category of URLs in Messages, page 15-7](#)

## Customizing the Appearance of End User Notification Page

Based on the evaluation by the Cisco Cloud Web Security proxy service, if the site is malicious, the end user sees a notice that the site is malicious and access to it has been blocked. You can customize the appearance of this notification page and display your organization's branding such as company logo, contact information, and so on.

**Note**

If you do not customize the notification page, end user sees the Cisco branded notification page.

**Before You Begin**

- Review the characteristics of the [Cisco Web Security Proxy End User Notification Page, page 15-5](#).
- Enable URL filtering. See [Enable URL Filtering, page 15-2](#).

**Procedure**

- 
- Step 1** Select **Security Services > Block Page Customization**.
- Step 2** Click **Enable**.
- Step 3** Check the **Enable Block Page customization** check box and enter the following details:
- URL of the organization's logo. It is recommended that the logo image is hosted on a publicly accessible server.
  - Organization's name
  - Organization's contact information
- Step 4** Choose the language of the notification. You can choose any one of the languages supported by the web interface.
- 
-  **Note** The default language of the end user's browser takes precedence over the language you have selected here. Also, if the default language of the end user's browser is not supported by AsyncOS, then the notification is displayed in the language you have selected here.
- 
- Step 5** (Optional) Preview the notification page by clicking **Preview Block Page Customization** link.
- Step 6** Submit and commit your changes.
- 

**Next Steps**

Set up URL rewriting in one of the following ways:

- Using Outbreak Filters. See [Redirecting URLs, page 14-5](#).
- Using Content or Message Filters. See [Taking Action Based on the Reputation or Category of URLs in Messages, page 15-7](#).

# Taking Action Based on the Reputation or Category of URLs in Messages

You can take action based on the reputation or category of URL links in messages using message filters and content filters in incoming and outgoing mail policies.

Because Outbreak Filters take many factors into consideration when evaluating messages for malware, and URL reputation alone may not trigger aggressive message handling, you may want to create filters based on URL reputation.

For example, you can use URL Reputation filters to:

- Rewrite URLs of suspect or unknown reputation to redirect them to the Cisco cloud Web Security proxy service for click-time evaluation.
- Drop messages that include URLs that have reputation scores in the Malicious range.

You can use URL Category filters to:

- Filter categories of URLs to enforce organizational policies for acceptable web use, for example to prevent users from visiting adult or gambling sites while at the office.
- Provide enhanced protection from malicious sites, which may not exist long enough to be classified. You can redirect all URLs in the Unclassified category to the Cisco cloud Web Security proxy service for evaluation at the time a user clicks a link.

## Related Topics

- [Using URL-Related Conditions \(Rules\) and Actions, page 15-7](#)
- [Filtering by URL Reputation or URL Category: Conditions and Rules, page 15-8](#)
- [Modifying URLs in Messages: Using URL Reputation and URL Category Actions in Filters, page 15-8](#)
- [Redirected URLs: What Does the End User Experience?, page 15-10](#)

## Using URL-Related Conditions (Rules) and Actions

| To                                                  | Example                                                                     | Do This                                                                                                                                                                                                                            |
|-----------------------------------------------------|-----------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Take action on the message as a whole.              | Drop or quarantine messages.                                                | Create a URL Reputation or URL Category condition or rule, then pair it with any action other than a URL Reputation or URL Category action.<br><br>Exception: Do not pair a URL Reputation condition or rule with a Bounce action. |
| Modify URLs in a message, or modify their behavior. | Replace a URL in the message with a text note, or make the URL unclickable. | Create a URL Reputation or URL Category action only; do not use a separate URL filtering condition.                                                                                                                                |

As always, you must specify a content filter in a mail policy in order to use it.

## Filtering by URL Reputation or URL Category: Conditions and Rules

You can perform actions on messages based on the reputation or category of URLs in the message. If you want to perform any action other than modifying URLs or their behavior, add a **URL Reputation** or **URL Category** condition and select the reputation scores or URL categories for which you want to apply the action.

For example, if you want to apply the **Drop (Final Action)** action to all messages that include URLs in the Adult category, add a condition of type **URL Category** with the **Adult** category selected.

If you do not specify a category, the action you choose is applied to all messages.

URL reputation score ranges for clean, suspect, and malicious URLs are predefined and not editable. However, you can specify a custom range instead. The specified endpoints are included in the range you specify. For example, if you create a custom range from -8 to -10, then -8 and -10 are included in the range. Use "No Score" for URLs for which a reputation score cannot be determined.

URLs that are included on the selected URL whitelist or on the global URL whitelist not evaluated.

The action that you pair with this condition is taken if any URL in the message matches the reputation score or any category specified in the condition.

If you want to modify URLs in a message, or modify their behavior, configure only a URL Reputation or URL Category action. You do not need a separate URL Reputation or URL Category condition or rule for this purpose.



### Note

Do not pair a URL Reputation condition with a Bounce action.



### Tip

To check the category of a particular URL, visit the link in [Reporting Uncategorized and Misclassified URLs, page 15-21](#).

### Related Topics

- [Chapter 11, “Content Filters”](#)
- [URL Reputation Rules, page 9-47](#)
- [URL Category Rule, page 9-47](#)
- [Creating Whitelists for URL Filtering, page 15-4](#)

## Modifying URLs in Messages: Using URL Reputation and URL Category Actions in Filters

Use a URL Reputation or URL Category action to modify URLs in a message, or their behavior, based on the reputation or category of the URL.

URL Reputation and URL Category actions do not require a separate condition. Instead, the selected action is applied based on the reputation or categories that you select in the URL Reputation or URL Category action.

The action is applied only to URLs that meet the criteria specified in the action. Other URLs in the message are not modified.

If you do not specify a category, the action you choose is applied to all messages.

URL reputation score ranges for clean, suspect, and malicious URLs are predefined and not editable. However, you can specify a custom range instead. The specified endpoints are included in the range you specify. For example, if you create a custom range from -8 to -10, then -8 and -10 are included in the range. Use "No Score" for URLs for which a reputation score cannot be determined.

The following URL-related actions are available:

- Defang a URL so that it is unclickable. Message recipients can still see and copy the URL.
- Redirect a URL so that if the message recipient clicks the link, the transaction is routed to a Cisco web security proxy in the cloud, which blocks access if the site is malicious.

Example: You might want to redirect all URLs in the **Uncategorized** category to the Cisco Cloud Web Security proxy service, as malicious sites used in phishing attacks often do not exist long enough to be classified.

See also [Redirected URLs: What Does the End User Experience?, page 15-10](#).

To redirect URLs to a different proxy, see the example in the following bullet.



**Note** The Cisco Cloud Web Security proxy service has no configurable options in this release. For example, there is no threat score threshold to adjust or action to specify based on threat score.

- Replace the URL with any text.

To include the original URL in the text that appears in the message, use the `$URL` variable.

Examples:

- Replace all URLs in the **Illegal Downloads** category with a note:

Message from your system administrator: A link to an illegal downloads web site has been removed from this message.

- Include the original URL along with a warning:

WARNING! The following URL may contain malware: `$URL`

This becomes: WARNING: The following URL may contain malware: `http://example.com`.

- Redirect to a custom proxy or web security service:

`http://custom_proxy/$URL`

This becomes: `http://custom_proxy/http://example.com`

The reputation and category of URLs that are included on the selected URL whitelist or on the global URL whitelist are not evaluated.

If you defang or replace URLs, you can choose to ignore URLs in signed messages.

Pairing a URL Reputation or URL Category action with a URL Reputation or URL Category condition (or rule) is not recommended. If you pair a condition (rule) and action that include different categories, then no match occurs.



**Tip**

To check the category of a particular URL, visit the link in [Reporting Uncategorized and Misclassified URLs, page 15-21](#).



**Related Topics**

- [Using Custom Headers to Redirect URLs in Suspected Spam to the Cisco Web Security Proxy: Configuration Example, page 13-11](#)
- [Chapter 11, “Content Filters”](#)
- [URL Reputation Actions, page 9-75](#)
- [URL Category Actions, page 9-76](#)
- [Creating Whitelists for URL Filtering, page 15-4](#)

## Redirected URLs: What Does the End User Experience?

Based on the evaluation by the Cisco Cloud Web Security proxy service:

- If the site is benign, the user is directed to the target web site and has no knowledge that the link has been redirected.
- If the site is malicious, the user sees a notice that the site is malicious and access to it has been blocked.

You can customize the appearance of end user notification page and display your organization's branding such as company logo, contact information, and so on. See [Customizing the Appearance of End User Notification Page, page 15-6](#).

- If communication with the Cisco Cloud Web Security proxy service times out, the user is allowed to access the target web site.
- If any other error occurs, the user sees a notice.

## Monitoring URL Filtering Results

To view data about malicious and suspicious URLs detected, select Monitor > URL Filtering. For important information about the data on this page, see [URL Filtering Page, page 28-21](#).

## Troubleshooting URL Filtering

**Related Topics**

- [Viewing Logs, page 15-11](#)
- [Alert: SDS: Error Fetching Enrollment Certificate, page 15-11](#)
- [Alert: SDS: Certificate Is Invalid, page 15-11](#)
- [Unable to Connect to Cisco Web Security Services, page 15-11](#)
- [Using the websecurityadvancedconfig Command, page 15-12](#)
- [Message Tracking Search Does Not Find Messages with Specified Category, page 15-12](#)
- [Malicious URLs and Marketing Messages Are Not Caught by Anti-Spam or Outbreak Filters, page 15-12](#)
- [URLs in a Filtered Category Are Not Handled Correctly, page 15-13](#)
- [End User Reaches Malicious Site via Rewritten URL, page 15-13](#)

- [Manually Configuring a Certificate for Communication with Cisco Web Security Services, page 15-13](#)

## Viewing Logs

URL filtering information is posted to the following logs:

- Mail Logs (`mail_logs`). Information related to the result of scanning a URL (action taken of a message depending on the URL) is posted to this log.
- URL Filtering Logs (`web_client`). Information related to errors, timeouts, network issues, and so on while attempting the URL lookup are posted this log.

Most information is at Info or Debug level.

Logs do not include information about what happens when a user clicks a redirected link in a message.

"SDS" in logs refers to URL reputation services.

## Alert: SDS: Error Fetching Enrollment Certificate

**Problem** You receive an info-level alert about an error fetching the enrollment client certificate.

**Solution** This certificate is required to connect to Cisco Web Security Services in the cloud in order to obtain URL reputation and category. Try the following:

- Check for networking issues such as incorrect proxy settings or firewall issues.
- Verify that your URL Filtering feature key is valid and active.
- If the problem persists, contact Cisco TAC.

## Alert: SDS: Certificate Is Invalid

**Problem** You receive a critical alert about an invalid SDS certificate.

**Solution** This certificate is required to connect to Cisco Web Security Services in the cloud in order to obtain URL reputation and category.

To obtain and manually install a certificate, see [Manually Configuring a Certificate for Communication with Cisco Web Security Services, page 15-13](#).

## Unable to Connect to Cisco Web Security Services

**Problem** The **Security Services > URL Filtering** page persistently indicates an issue connecting to Cisco Web Security Services.

**Solution**

- If you have enabled URL filtering but have not yet committed the change, commit the change.
- Check for recent alerts related to the connection with Cisco Web Security Services. See [Viewing Recent Alerts, page 33-38](#). If applicable, see [Alert: SDS: Error Fetching Enrollment Certificate, page 15-11](#) and [Alert: SDS: Certificate Is Invalid, page 15-11](#).

- If you are connecting via a proxy specified in Security Services > Service Updates, verify that this is configured and working properly.
- Check for other network issues that might prevent connection.
- If you see errors in the URL Filtering Logs related to timed out requests to the SDS client, use the `websecuritydiagnostics` command and the `websecurityadvancedconfig` command in the command-line interface to investigate and make changes:
  - If the diagnostics show that Response Time or DNS Lookup Time is not less than the configured URL Lookup Timeout, increase the URL Lookup Timeout value accordingly.
  - If the diagnostics show that the cache size is at or near the capacity specified in the advanced configuration settings, increase the cache size.
- Check the URL Filtering Logs for non-timeout errors in communications with the URL scanner, Cisco Web Security Services, or SDS. "SDS" in logs represents Cisco Web Security Services. If you see such log messages, contact TAC.

## Using the `websecurityadvancedconfig` Command

Except for changes explicitly described in this document, make no other changes using the `websecurityadvancedconfig` command without guidance from TAC.

## Message Tracking Search Does Not Find Messages with Specified Category

**Problem** Messages that contain URLs in a particular category are not found when searching by that category.

**Solution** See [Expected Messages Are Missing from Search Results, page 29-7](#).

## Malicious URLs and Marketing Messages Are Not Caught by Anti-Spam or Outbreak Filters

**Problem** Malicious URLs and messages containing marketing links are not caught by the anti-spam or outbreak filters.

### Solution

- This can occur because web site reputation and category are only two among many criteria that anti-spam and outbreak filters use to determine their verdicts. You can increase the sensitivity of these filters by lowering the thresholds required to take action such as rewriting or replacing URLs with text, or quarantining or dropping messages. For details, see [The Outbreak Filters Feature and Mail Policies, page 14-15](#) and [Defining Anti-Spam Policies, page 13-7](#). Alternatively, create content or message filters based on URL reputation score.
- This can also occur if the Email Security appliance is unable to connect to the Cisco Web Security Services. See [Unable to Connect to Cisco Web Security Services, page 15-11](#).

## URLs in a Filtered Category Are Not Handled Correctly

**Problem** The defined action in a content or message filter based on URL category is not applied.

**Solution**

- Use the Trace feature (described in the Troubleshooting chapter) to follow the message processing path.
- This can occur if the Email Security appliance is unable to connect to the Cisco Web Security Services. See [Unable to Connect to Cisco Web Security Services, page 15-11](#).
- If there are no connection issues, the URLs may not yet be categorized, or may be miscategorized. See [Reporting Uncategorized and Misclassified URLs, page 15-21](#). You can use this site to determine the category of a URL.

## End User Reaches Malicious Site via Rewritten URL

**Problem** A malicious URL was redirected to the Cisco Web Security Proxy, but the end user was able to access the site anyway.

**Solution** This can occur if:

- The site was not yet identified as a malicious site.
- The connection to the Cisco Web Security Proxy timed out, which should be a rare occurrence. Ensure that network issues are not interfering with the connection.

## Manually Configuring a Certificate for Communication with Cisco Web Security Services

Use this procedure if the appliance is unable to automatically obtain a certificate for communication with Cisco Web Security Services.

**Procedure**

- 
- |               |                                                                                                                                           |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Obtain the required certificate.                                                                                                          |
| <b>Step 2</b> | Upload the certificate using <b>Network &gt; Certificates</b> , or use the <code>certconfig</code> command in the command-line interface. |
| <b>Step 3</b> | In the command-line interface, enter the <code>websecurityconfig</code> command.                                                          |
| <b>Step 4</b> | Follow the prompts to set the client certificate for Cisco Web Security Services Authentication.                                          |
| <b>Step 5</b> | After completing the certificate installation process, enter the <code>webcacheflush</code> command.                                      |
- 

## About URL Categories

**Related Topics**

- [URL Category Descriptions, page 15-14](#)

- [Determining the Category of a URL, page 15-21](#)
- [Reporting Uncategorized and Misclassified URLs, page 15-21](#)
- [Future URL Category Set Changes, page 15-22](#)

## URL Category Descriptions

These URL categories are the same categories that are used on recent releases of AsyncOS for Web Security appliances.

**Table 15-1**

| URL Category   | Abbreviation | Code | Description                                                                                                                                                                                                                                                                                                            | Example URLs                                           |
|----------------|--------------|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------|
| Adult          | adlt         | 1006 | Directed at adults, but not necessarily pornographic. May include adult clubs (strip clubs, swingers clubs, escort services, strippers); general information about sex, non-pornographic in nature; genital piercing; adult products or greeting cards; information about sex not in the context of health or disease. | www.adultentertainmentexpo.com<br>www.adultnetline.com |
| Advertisements | adv          | 1027 | Banner and pop-up advertisements that often accompany a web page; other advertising websites that provide advertisement content. Advertising services and sales are classified as “Business and Industry.”                                                                                                             | www.adforce.com<br>www.doubleclick.com                 |
| Alcohol        | alc          | 1077 | Alcohol as a pleasurable activity; beer and wine making, cocktail recipes; liquor sellers, wineries, vineyards, breweries, alcohol distributors. Alcohol addiction is classified as “Health and Nutrition.” Bars and restaurants are classified as “Dining and Drinking.”                                              | www.samueladams.com<br>www.whisky.com                  |
| Arts           | art          | 1002 | Galleries and exhibitions; artists and art; photography; literature and books; performing arts and theater; musicals; ballet; museums; design; architecture. Cinema and television are classified as “Entertainment.”                                                                                                  | www.moma.org<br>www.nga.gov                            |
| Astrology      | astr         | 1074 | Astrology; horoscope; fortune telling; numerology; psychic advice; tarot.                                                                                                                                                                                                                                              | www.astro.com<br>www.astrology.com                     |
| Auctions       | auct         | 1088 | Online and offline auctions, auction houses, and classified advertisements.                                                                                                                                                                                                                                            | www.craigslist.com<br>www.ebay.com                     |

Table 15-1

| URL Category               | Abbreviation | Code | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Example URLs                                                       |
|----------------------------|--------------|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------|
| Business and Industry      | busi         | 1019 | Marketing, commerce, corporations, business practices, workforce, human resources, transportation, payroll, security and venture capital; office supplies; industrial equipment (process equipment), machines and mechanical systems; heating equipment, cooling equipment; materials handling equipment; packaging equipment; manufacturing: solids handling, metal fabrication, construction and building; passenger transportation; commerce; industrial design; construction, building materials; shipping and freight (freight services, trucking, freight forwarders, truckload carriers, freight and transportation brokers, expedited services, load and freight matching, track and trace, rail shipping, ocean shipping, road feeder services, moving and storage). | www.freightcenter.com<br>www.staples.com                           |
| Chat and Instant Messaging | chat         | 1040 | Web-based instant messaging and chat rooms.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | www.icq.com<br>www.meebo.com                                       |
| Cheating and Plagiarism    | plag         | 1051 | Promoting cheating and selling written work, such as term papers, for plagiarism.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | www.bestessays.com<br>www.superiorpapers.com                       |
| Child Abuse Content        | cprn         | 1064 | Worldwide illegal child sexual abuse content.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | —                                                                  |
| Computer Security          | csec         | 1065 | Offering security products and services for corporate and home users.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | www.computersecurity.com<br>www.symantec.com                       |
| Computers and Internet     | comp         | 1003 | Information about computers and software, such as hardware, software, software support; information for software engineers, programming and networking; website design; the web and Internet in general; computer science; computer graphics and clip art. “Freeware and Shareware” is a separate category.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | www.xml.com<br>www.w3.org                                          |
| Dating                     | date         | 1055 | Dating, online personals, matrimonial agencies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | www.eharmony.com<br>www.match.com                                  |
| Digital Postcards          | card         | 1082 | Enabling sending of digital postcards and e-cards.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | www.all-yours.net<br>www.delivr.net                                |
| Dining and Drinking        | food         | 1061 | Eating and drinking establishments; restaurants, bars, taverns, and pubs; restaurant guides and reviews.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | www.hideawaybrewpub.com<br>www.restaurantrow.com                   |
| Dynamic and Residential    | dyn          | 1091 | IP addresses of broadband links that usually indicates users attempting to access their home network, for example for a remote session to a home computer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | http://109.60.192.55<br>http://dynalink.co.jp<br>http://ipadsl.net |

Table 15-1

| URL Category           | Abbreviation | Code | Description                                                                                                                                                                                                                                                                                                                                | Example URLs                                        |
|------------------------|--------------|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------|
| Education              | edu          | 1001 | Education-related, such as schools, colleges, universities, teaching materials, and teachers' resources; technical and vocational training; online training; education issues and policies; financial aid; school funding; standards and testing.                                                                                          | www.education.com<br>www.greatschools.org           |
| Entertainment          | ent          | 1093 | Details or discussion of films; music and bands; television; celebrities and fan websites; entertainment news; celebrity gossip; entertainment venues. Compare with the "Arts" category.                                                                                                                                                   | www.eonline.com<br>www.ew.com                       |
| Extreme                | extr         | 1075 | Material of a sexually violent or criminal nature; violence and violent behavior; tasteless, often gory photographs, such as autopsy photos; photos of crime scenes, crime and accident victims; excessive obscene material; shock websites.                                                                                               | www.car-accidents.com<br>www.crime-scene-photos.com |
| Fashion                | fash         | 1076 | Clothing and fashion; hair salons; cosmetics; accessories; jewelry; perfume; pictures and text relating to body modification; tattoos and piercing; modeling agencies. Dermatological products are classified as "Health and Nutrition."                                                                                                   | www.fashion.net<br>www.findabeautysalon.com         |
| File Transfer Services | fts          | 1071 | File transfer services with the primary purpose of providing download services and hosted file sharing                                                                                                                                                                                                                                     | www.rapidshare.com<br>www.yousendit.com             |
| Filter Avoidance       | filt         | 1025 | Promoting and aiding undetectable and anonymous web usage, including cgi, php and glype anonymous proxy services.                                                                                                                                                                                                                          | www.bypassschoolfilter.com<br>www.filterbypass.com  |
| Finance                | fnnc         | 1015 | Primarily financial in nature, such as accounting practices and accountants, taxation, taxes, banking, insurance, investing, the national economy, personal finance involving insurance of all types, credit cards, retirement and estate planning, loans, mortgages. Stock and shares are classified as "Online Trading."                 | finance.yahoo.com<br>www.bankofamerica.com          |
| Freeware and Shareware | free         | 1068 | Providing downloads of free and shareware software.                                                                                                                                                                                                                                                                                        | www.freewarehome.com<br>www.shareware.com           |
| Gambling               | gamb         | 1049 | Casinos and online gambling; bookmakers and odds; gambling advice; competitive racing in a gambling context; sports booking; sports gambling; services for spread betting on stocks and shares. Websites dealing with gambling addiction are classified as "Health and Nutrition." Government-run lotteries are classified as "Lotteries". | www.888.com<br>www.gambling.com                     |
| Games                  | game         | 1007 | Various card games, board games, word games, and video games; combat games; sports games; downloadable games; game reviews; cheat sheets; computer games and Internet games, such as role-playing games.                                                                                                                                   | www.games.com<br>www.shockwave.com                  |

Table 15-1

| URL Category         | Abbreviation | Code | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Example URLs                              |
|----------------------|--------------|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|
| Government and Law   | gov          | 1011 | Government websites; foreign relations; news and information relating to government and elections; information relating to the field of law, such as attorneys, law firms, law publications, legal reference material, courts, dockets, and legal associations; legislation and court decisions; civil rights issues; immigration; patents and copyrights; information relating to law enforcement and correctional systems; crime reporting, law enforcement, and crime statistics; military, such as the armed forces, military bases, military organizations; anti-terrorism. | www.usa.gov<br>www.law.com                |
| Hacking              | hack         | 1050 | Discussing ways to bypass the security of websites, software, and computers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | www.hackthissite.org<br>www.gohacking.com |
| Hate Speech          | hate         | 1016 | Websites promoting hatred, intolerance, or discrimination on the basis of social group, color, religion, sexual orientation, disability, class, ethnicity, nationality, age, gender, gender identity; sites promoting racism; sexism; racist theology; hate music; neo-Nazi organizations; supremacism; Holocaust denial.                                                                                                                                                                                                                                                        | www.kkk.com<br>www.nazi.org               |
| Health and Nutrition | hlth         | 1009 | Health care; diseases and disabilities; medical care; hospitals; doctors; medicinal drugs; mental health; psychiatry; pharmacology; exercise and fitness; physical disabilities; vitamins and supplements; sex in the context of health (disease and health care); tobacco use, alcohol use, drug use, and gambling in the context of health (disease and health care); food in general; food and beverage; cooking and recipes; food and nutrition, health, and dieting; cooking, including recipe and culinary websites; alternative medicine.                                 | www.health.com<br>www.webmd.com           |
| Humor                | lol          | 1079 | Jokes, sketches, comics and other humorous content. Adult humor likely to offend is classified as “Adult.”                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | www.humor.com<br>www.jokes.com            |
| Illegal Activities   | ilac         | 1022 | Promoting crime, such as stealing, fraud, illegally accessing telephone networks; computer viruses; terrorism, bombs, and anarchy; websites depicting murder and suicide as well as explaining ways to commit them.                                                                                                                                                                                                                                                                                                                                                              | www.ekran.no<br>www.thedisease.net        |
| Illegal Downloads    | ildl         | 1084 | Providing the ability to download software or other materials, serial numbers, key generators, and tools for bypassing software protection in violation of copyright agreements. Torrents are classified as “Peer File Transfer.”                                                                                                                                                                                                                                                                                                                                                | www.keygenguru.com<br>www.zcrack.com      |



Table 15-1

| URL Category                                 | Abbreviation | Code | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Example URLs                                 |
|----------------------------------------------|--------------|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| Illegal Drugs                                | drug         | 1047 | Information about recreational drugs, drug paraphernalia, drug purchase and manufacture.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | www.cocaine.org<br>www.hightimes.com         |
| Infrastructure and Content Delivery Networks | infr         | 1018 | Content delivery infrastructure and dynamically generated content; websites that cannot be classified more specifically because they are secured or otherwise difficult to classify.                                                                                                                                                                                                                                                                                                                                                                                                                                                             | www.akamai.net<br>www.webstat.net            |
| Internet Telephony                           | voip         | 1067 | Telephonic services using the Internet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | www.evaphone.com<br>www.skype.com            |
| Job Search                                   | job          | 1004 | Career advice; resume writing and interviewing skills; job placement services; job databanks; permanent and temporary employment agencies; employer websites.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | www.careerbuilder.com<br>www.monster.com     |
| Lingerie and Swimsuits                       | ling         | 1031 | Intimate apparel and swim wear, especially when modeled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | www.swimsuits.com<br>www.victoriassecret.com |
| Lotteries                                    | lotr         | 1034 | Sweepstakes, contests and state-sponsored lotteries.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | www.calottery.com<br>www.flalottery.com      |
| Mobile Phones                                | cell         | 1070 | Short Message Services (SMS); ring tones and mobile phone downloads. Cellular carrier websites are included in the “Business and Industry” category.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | www.cbfsms.com<br>www.zedge.net              |
| Nature                                       | natr         | 1013 | Natural resources; ecology and conservation; forests; wilderness; plants; flowers; forest conservation; forest, wilderness, and forestry practices; forest management (reforestation, forest protection, conservation, harvesting, forest health, thinning, and prescribed burning); agricultural practices (agriculture, gardening, horticulture, landscaping, planting, weed control, irrigation, pruning, and harvesting); pollution issues (air quality, hazardous waste, pollution prevention, recycling, waste management, water quality, and the environmental cleanup industry); animals, pets, livestock, and zoology; biology; botany. | www.enature.com<br>www.nature.org            |
| News                                         | news         | 1058 | News; headlines; newspapers; television stations; magazines; weather; ski conditions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | www.cnn.com<br>news.bbc.co.uk                |
| Non-Governmental Organizations               | ngo          | 1087 | Non-governmental organizations such as clubs, lobbies, communities, non-profit organizations and labor unions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | www.panda.org<br>www.unions.org              |
| Non-Sexual Nudity                            | nsn          | 1060 | Nudism and nudity; naturism; nudist camps; artistic nudes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | www.artenuda.com<br>www.naturistsociety.com  |

Table 15-1

| URL Category              | Abbreviation | Code | Description                                                                                                                                                                                                                                                                                                                                                                   | Example URLs                              |
|---------------------------|--------------|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|
| Online Communities        | comm         | 1024 | Affinity groups; special interest groups; web newsgroups; message boards. Excludes websites classified as “Professional Networking” or “Social Networking.”                                                                                                                                                                                                                   | www.igda.org<br>www.ieee.org              |
| Online Storage and Backup | osb          | 1066 | Offsite and peer-to-peer storage for backup, sharing, and hosting.                                                                                                                                                                                                                                                                                                            | www.adrive.com<br>www.dropbox.com         |
| Online Trading            | trad         | 1028 | Online brokerages; websites that enable the user to trade stocks online; information relating to the stock market, stocks, bonds, mutual funds, brokers, stock analysis and commentary, stock screens, stock charts, IPOs, stock splits. Services for spread betting on stocks and shares are classified as “Gambling.” Other financial services are classified as “Finance.” | www.tdameritrade.com<br>www.scottrade.com |
| Organizational Email      | pem          | 1085 | Websites used to access business email (often via Outlook Web Access).                                                                                                                                                                                                                                                                                                        | —                                         |
| Parked Domains            | park         | 1092 | Websites that monetize traffic from the domain using paid listings from an ad network, or are owned by “squatters” hoping to sell the domain name for a profit. These also include fake search websites which return paid ad links.                                                                                                                                           | www.domainzaar.com<br>www.parked.com      |
| Peer File Transfer        | p2p          | 1056 | Peer-to-peer file request websites. This does not track the file transfers themselves.                                                                                                                                                                                                                                                                                        | www.bittorrent.com<br>www.limewire.com    |
| Personal Sites            | pers         | 1081 | Websites about and from private individuals; personal homepage servers; websites with personal contents; personal blogs with no particular theme.                                                                                                                                                                                                                             | www.karymullis.com<br>www.stallman.org    |
| Photo Searches and Images | img          | 1090 | Facilitating the storing and searching for, images, photographs, and clip-art.                                                                                                                                                                                                                                                                                                | www.flickr.com<br>www.photobucket.com     |
| Politics                  | pol          | 1083 | Websites of politicians; political parties; news and information on politics, elections, democracy, and voting.                                                                                                                                                                                                                                                               | www.politics.com<br>www.thisnation.com    |
| Pornography               | porn         | 1054 | Sexually explicit text or depictions. Includes explicit anime and cartoons; general explicit depictions; other fetish material; explicit chat rooms; sex simulators; strip poker; adult movies; lewd art; web-based explicit email.                                                                                                                                           | www.redtube.com<br>www.youporn.com        |
| Professional Networking   | pnet         | 1089 | Social networking for the purpose of career or professional development. See also “Social Networking.”                                                                                                                                                                                                                                                                        | www.linkedin.com<br>www.europeanpwn.net   |
| Real Estate               | rest         | 1045 | Information that would support the search for real estate; office and commercial space; real estate listings, such as rentals, apartments, and homes; house building.                                                                                                                                                                                                         | www.realtor.com<br>www.zillow.com         |

Table 15-1

| URL Category               | Abbreviation | Code | Description                                                                                                                                                                                                                                                 | Example URLs                                        |
|----------------------------|--------------|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------|
| Reference                  | ref          | 1017 | City and state guides; maps, time; reference sources; dictionaries; libraries.                                                                                                                                                                              | www.wikipedia.org<br>www.yellowpages.com            |
| Religion                   | rel          | 1086 | Religious content, information about religions; religious communities.                                                                                                                                                                                      | www.religionfacts.com<br>www.religioustolerance.org |
| SaaS and B2B               | saas         | 1080 | Web portals for online business services; online meetings.                                                                                                                                                                                                  | www.netsuite.com<br>www.salesforce.com              |
| Safe for Kids              | kids         | 1057 | Directed at, and specifically approved for, young children.                                                                                                                                                                                                 | kids.discovery.com<br>www.nickjr.com                |
| Science and Technology     | sci          | 1012 | Science and technology, such as aerospace, electronics, engineering, mathematics, and other similar subjects; space exploration; meteorology; geography; environment; energy (fossil, nuclear, renewable); communications (telephones, telecommunications). | www.physorg.com<br>www.science.gov                  |
| Search Engines and Portals | srch         | 1020 | Search engines and other initial points of access to information on the Internet.                                                                                                                                                                           | www.bing.com<br>www.google.com                      |
| Sex Education              | sxed         | 1052 | Factual websites dealing with sex; sexual health; contraception; pregnancy.                                                                                                                                                                                 | www.avert.org<br>www.scarleteen.com                 |
| Shopping                   | shop         | 1005 | Bartering; online purchasing; coupons and free offers; general office supplies; online catalogs; online malls.                                                                                                                                              | www.amazon.com<br>www.shopping.com                  |
| Social Networking          | snet         | 1069 | Social networking. See also “Professional Networking.”                                                                                                                                                                                                      | www.facebook.com<br>www.twitter.com                 |
| Social Science             | socs         | 1014 | Sciences and history related to society; archaeology; anthropology; cultural studies; history; linguistics; geography; philosophy; psychology; women's studies.                                                                                             | www.archaeology.org<br>www.anthropology.net         |
| Society and Culture        | scty         | 1010 | Family and relationships; ethnicity; social organizations; genealogy; seniors; child-care.                                                                                                                                                                  | www.childcare.gov<br>www.familysearch.org           |
| Software Updates           | swup         | 1053 | Websites that host updates for software packages.                                                                                                                                                                                                           | www.softwarepatch.com<br>www.versiontracker.com     |
| Sports and Recreation      | sprt         | 1008 | All sports, professional and amateur; recreational activities; fishing; fantasy sports; public parks; amusement parks; water parks; theme parks; zoos and aquariums; spas.                                                                                  | www.espn.com<br>www.recreation.gov                  |
| Streaming Audio            | aud          | 1073 | Real-time streaming audio content including Internet radio and audio feeds.                                                                                                                                                                                 | www.live-radio.net<br>www.shoutcast.com             |
| Streaming Video            | vid          | 1072 | Real-time streaming video including Internet television, web casts, and video sharing.                                                                                                                                                                      | www.hulu.com<br>www.youtube.com                     |

Table 15-1

| URL Category         | Abbreviation | Code | Description                                                                                                                                                                                                                                                                                                                            | Example URLs                                |
|----------------------|--------------|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------|
| Tobacco              | tob          | 1078 | Pro-tobacco websites; tobacco manufacturers; pipes and smoking products (not marketed for illegal drug use). Tobacco addiction is classified as “Health and Nutrition.”                                                                                                                                                                | www.bat.com<br>www.tobacco.org              |
| Transportation       | trns         | 1044 | Personal transportation; information about cars and motorcycles; shopping for new and used cars and motorcycles; car clubs; boats, airplanes, recreational vehicles (RVs), and other similar items. Note, car and motorcycle racing is classified as “Sports and Recreation.”                                                          | www.cars.com<br>www.motorcycles.com         |
| Travel               | trvl         | 1046 | Business and personal travel; travel information; travel resources; travel agents; vacation packages; cruises; lodging and accommodation; travel transportation; flight booking; airfares; car rental; vacation homes.                                                                                                                 | www.expedia.com<br>www.lonelyplanet.com     |
| Unclassified         | —            | —    | Websites which are not in the Cisco database are recorded as unclassified for reporting purposes. This may include mistyped URLs.                                                                                                                                                                                                      | —                                           |
| Weapons              | weap         | 1036 | Information relating to the purchase or use of conventional weapons such as gun sellers, gun auctions, gun classified ads, gun accessories, gun shows, and gun training; general information about guns; other weapons and graphic hunting sites may be included. Government military websites are classified as “Government and Law.” | www.coldsteel.com<br>www.gunbroker.com      |
| Web Hosting          | whst         | 1037 | Website hosting; bandwidth services.                                                                                                                                                                                                                                                                                                   | www.bluehost.com<br>www.godaddy.com         |
| Web Page Translation | tran         | 1063 | Translation of web pages between languages.                                                                                                                                                                                                                                                                                            | babelfish.yahoo.com<br>translate.google.com |
| Web-Based Email      | mail         | 1038 | Public web-based email services. Websites enabling individuals to access their company or organization’s email service are classified as “Organizational Email.”                                                                                                                                                                       | mail.yahoo.com<br>www.hotmail.com           |

## Determining the Category of a URL

To look up the category of a particular URL, visit the site shown in [Reporting Uncategorized and Misclassified URLs, page 15-21](#).

## Reporting Uncategorized and Misclassified URLs

To report URLs that have been miscategorized, and URLs that are not categorized but should be, visit:

[https://securityhub.cisco.com/web/submit\\_urls](https://securityhub.cisco.com/web/submit_urls)

To check the status of submitted URLs, click the **Status on Submitted URLs** tab on this page.

## Future URL Category Set Changes

Rarely, the set of URL categories may change as a result of emerging trends and technologies. For example, a category may be added or removed, renamed, merged with another category, or split into two categories. These changes can affect the results from existing filters, so if changes occur, the appliance will send an alert (System type, Warning severity). If you receive such an alert, you should evaluate and possibly update content and message filters to work with the updated categories. Existing filters will not automatically be changed. To ensure that you receive these alerts, see [Adding Alert Recipients, page 33-36](#).

The following changes do not require category set changes and will not generate alerts:

- Routine categorization of newly-categorized sites.
- Recategorization of misclassified sites.



## File Reputation Filtering and File Analysis

---

- [Overview of File Reputation Filtering and File Analysis, page 16-1](#)
- [Configuring File Reputation and Analysis Features, page 16-4](#)
- [File Reputation and File Analysis Reporting and Tracking, page 16-11](#)
- [Taking Action When File Threat Verdicts Change, page 16-13](#)
- [Troubleshooting File Reputation and Analysis, page 16-13](#)

### Overview of File Reputation Filtering and File Analysis

Advanced Malware Protection uses cloud-based services to protect against zero-day and targeted file-based threats in email attachments by:

- Obtaining each file's reputation.
- Analyzing behavior of certain files with unknown reputations.
- Notifying you about files that are determined to be threats after they have entered your network.

These features are available only for incoming messages. Files attached to outgoing messages are not evaluated.

### File Threat Verdict Updates

Threat verdicts can change as new information emerges. A file may initially be evaluated as unknown or clean, and the file may therefore be released to the recipient. If the threat verdict changes, you will be alerted, and the file and its new verdict appear in the AMP Verdict Updates report. You can investigate the point-of-entry message as a starting point to remediating any impacts of the threat.

Verdicts can also change from malicious to clean.

When the appliance processes subsequent instances of the same file, the updated verdict is immediately applied.

#### Related Topics

- [File Reputation and File Analysis Reporting and Tracking, page 16-11](#)
- [Taking Action When File Threat Verdicts Change, page 16-13](#)


## File Processing Overview

Evaluation of file reputation and sending of files for analysis occur immediately after anti-virus scanning, regardless of verdicts from previous scanning engines, unless a final action has been taken on the message.

Communications between the appliance and the file reputation service are encrypted and protected from tampering.

After a file's reputation is evaluated:

- If the file is known to the file reputation service and is determined to be clean, the message continues through the workqueue.
- If the file reputation service returns a verdict of malicious for any attachment in the message, then the appliance applies the action that you have specified in the applicable mail policy.
- If the file is known to the reputation service but there is insufficient information for a definitive verdict, the reputation service returns a reputation score based on characteristics of the file such as threat fingerprint and behavioral analysis. If this score meets or exceeds the configured reputation threshold (you should not change the default threshold), the appliance applies the action that you have configured in the mail policy for files that contain malware.
- If the reputation service has no information about the file, and the file does not meet the criteria for analysis, the file is considered clean and the message continues through the workqueue.
- If the reputation service has no information about the file, and the file meets the criteria for files that can be analyzed (see [Which Files Are Evaluated and Analyzed?](#), page 16-3), then the file is considered clean and is optionally sent for analysis. You can configure the appliance to quarantine files sent for analysis instead of releasing them immediately to the workqueue. See [Quarantining Messages with Attachments Sent for Analysis](#), page 16-7.
- If file reputation information is unavailable, for example because the connection with the cloud service timed out, the appliance applies the action that you have specified for unscannable attachments in the applicable mail policy.



371477

If the file is sent for analysis:

- Files are sent using SSL/TLS.
- Analysis normally takes minutes, but may take longer.
- Information about every file analyzed is added to the reputation database. File analysis results contribute to the reputation of the file.

For information about verdict updates, see [File Threat Verdict Updates](#), page 16-1.

## Which Files Are Evaluated and Analyzed?

The reputation service evaluates most file types. File type identification is determined by file content and is not dependent on the filename extension.

Some files with unknown reputation can be analyzed for threat characteristics. When you configure the file analysis feature, you choose which file types are sent for analysis. New types can be added dynamically; you will receive an alert when the list of uploadable file types changes, and can select added file types to upload.

For complete information about which files are evaluated and analyzed, see the Release Notes for your AsyncOS version, available from:

<http://www.cisco.com/c/en/us/support/security/email-security-appliance/products-release-notes-list.html>.

### Related Topics

- [Archive or Compressed File Processing](#), page 16-4



## Archive or Compressed File Processing

If the file is compressed or archived,

- Reputation of the compressed or archive file is evaluated.
- The compressed or archive file is decompressed and reputations of all the extracted files are evaluated.

For information about which archived and compressed files are examined, including file formats, see *File Criteria for Advanced Malware Protection Services for Cisco Content Security Products*, available from

<http://www.cisco.com/c/en/us/support/security/email-security-appliance/products-user-guide-list.html>.

In this scenario,

- If one of the extracted files is malicious, the file reputation service returns a verdict of Malicious for the compressed or the archive file.
- If the compressed or archive file is malicious and all the extracted files are clean, the file reputation service returns a verdict of Malicious for the compressed or the archive file.
- If the verdict of any of the extracted files is unknown, the extracted files are optionally (if configured and the file type is supported for file analysis) sent for file analysis.
- If the extraction of a file fails while decompressing a compressed or an archive file, the file reputation service returns a verdict of Unscannable for the compressed or the archive file. Keep in mind that, in this scenario, if one of the extracted files is malicious, the file reputation service returns a verdict of Malicious for the compressed or the archive file (Malicious verdict takes precedence over Unscannable verdict).

**Note**

Reputation of the extracted files with safe MIME types, for example, text/plain, are not evaluated.

## FIPS Compliance

File reputation scanning and file analysis are FIPS compliant.

## Configuring File Reputation and Analysis Features

- [Requirements for Communication with File Reputation and Analysis Services, page 16-5](#)
- [Enabling and Configuring File Reputation and Analysis Services, page 16-5](#)
- [Configuring the Incoming Mail Policy for File Reputation Scanning and File Analysis, page 16-6](#)
- [Quarantining Messages with Attachments Sent for Analysis, page 16-7](#)
- [Using File Analysis Quarantine, page 16-8](#)
- [Centralized File Analysis Quarantine, page 16-10](#)
- [X-Headers for File Reputation and Analysis, page 16-10](#)
- [Sending Notifications to End Users about Dropped Messages or Attachments, page 16-10](#)
- [Advanced Malware Protection and Clusters, page 16-10](#)
- [Ensuring That You Receive Alerts, page 16-11](#)

- [Configuring Centralized Reporting for Advanced Malware Protection Features, page 16-11](#)

## Requirements for Communication with File Reputation and Analysis Services

- All Email Security appliances that use these services must be able to connect to them directly over the internet.
- By default, communication with file reputation and analysis services is routed through the interface that is associated with the default gateway. To route this traffic through a different interface, create a static route for each address in the Advanced section of the Security Services > File Reputation and Analysis page.
- Communication with cloud services for file reputation and analysis occurs over IPv4.
- For information about required open firewall ports, see [Appendix D, “Firewall Information.”](#)

### Related Topics

- [Configuring TCP/IP Traffic Routes, page 32-52](#)

## Enabling and Configuring File Reputation and Analysis Services

### Before You Begin

- Acquire feature keys for the file reputation service and the file analysis service.
- Meet the [Requirements for Communication with File Reputation and Analysis Services, page 16-5](#).

### Procedure

- 
- |               |                                                                                                                                                                 |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Select <b>Security Services &gt; File Reputation and Analysis</b> .                                                                                             |
| <b>Step 2</b> | Click <b>Enable</b> .                                                                                                                                           |
| <b>Step 3</b> | Click <b>Edit Global Settings</b> .                                                                                                                             |
| <b>Step 4</b> | Select <b>Enable File Reputation</b> .                                                                                                                          |
| <b>Step 5</b> | Accept the license agreement if presented.                                                                                                                      |
| <b>Step 6</b> | File Analysis is enabled by default. If you do not uncheck <b>Enable File Analysis</b> , the File Analysis feature key will be activated after the next commit. |
| <b>Step 7</b> | In the File Analysis section, select the file types to send to the cloud for analysis.                                                                          |

**Step 8** Adjust the following Advanced settings as desired:

| Option                                                                                                                              | Description                                                                                                                                                                                                                                                                                                                          |
|-------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SSL Communication for File Reputation                                                                                               | <p>Check <b>Use SSL (Port 443)</b> to communicate on port 443 instead of the default port, 32137.</p> <p>This option also allows you to configure an upstream proxy for communication with the file reputation service.</p> <p><b>Note</b> SSL communication over port 32137 may require you to open that port in your firewall.</p> |
| Routing Table                                                                                                                       | The routing table (associated with an appliance network interface type, either Management or Data) to be used for Advanced Malware Protection services. If the appliance has both the Management interface and one or more Data interfaces enabled, you can select Management or Data.                                               |
| Reputation Threshold <ul style="list-style-type: none"> <li>• Use value from Cloud Service</li> <li>• Enter custom value</li> </ul> | The upper limit for acceptable file reputation scores. Scores above this threshold indicate the file is infected.                                                                                                                                                                                                                    |



**Note** Do not change any other Advanced settings without guidance from Cisco support.

**Step 9** Submit and commit your changes.

## Configuring the Incoming Mail Policy for File Reputation Scanning and File Analysis

### Procedure

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

**Step 2** Click the link in the **Advanced Malware Protection** column of the mail policy to modify.

**Step 3** Choose options.

- If you do not want to send files to the cloud, for example for confidentiality reasons, uncheck **Enable File Analysis**.
- Select the actions that AsyncOS must perform if an attachment is considered Unscannable. Attachments are considered Unscannable when the appliance is unable to obtain information from the file reputation service for any reason, for example because the connection timed out.

Select the following:

- Whether to deliver or drop the message.
- Whether to archive the original message. Archived messages are stored as an mbox-format log file in the `amparchive` directory on the appliance. The preconfigured AMP Archive (`amparchive`) log subscription is required.

- Whether to warn the end user by modifying the message subject, for example, [WARNING: ATTACHMENT(S) MAY CONTAIN MALWARE].
- Whether to add a custom header to provide granular controls to the administrator.
- Select the actions that AsyncOS must perform if an attachment is considered Malicious. Select the following:
  - Whether to deliver or drop the message.
  - Whether to archive the original message. Archived messages are stored as an mbox-format log file in the `amparchive` directory on the appliance. The preconfigured AMP Archive (`amparchive`) log subscription is required.
  - Whether to deliver the message after removing the malware attachments.
  - Whether to warn the end user by modifying the message subject, for example, [WARNING: MALWARE DETECTED IN ATTACHMENT(S)].
  - Whether to add a custom header to provide granular controls to the administrator.
- Select the actions that AsyncOS must perform if an attachment is sent for File Analysis. Select the following:
  - Whether to deliver or quarantine the message.
  - Whether to archive the original message. Archived messages are stored as an mbox-format log file in the `amparchive` directory on the appliance. The preconfigured AMP Archive (`amparchive`) log subscription is required.
  - Whether to warn the end user by modifying the message subject, for example, “[WARNING: ATTACHMENT(S) MAY CONTAIN MALWARE].”
  - Whether to add a custom header to provide granular controls to the administrator.

**Step 4** Submit and commit your changes.

---

## Quarantining Messages with Attachments Sent for Analysis

You can configure the appliance to quarantine files sent for analysis instead of releasing them immediately to the workqueue. Quarantined messages and their attachments are rescanned for threats upon release from quarantine. If the message is released after file analysis results are available to the reputation scanner, any identified threats will be caught during rescanning.

### Procedure

---

- Step 1** Select **Mail Policies > Incoming Mail Policies**.
- Step 2** Click the link in the **Advanced Malware Protection** column of the mail policy to modify.
- Step 3** Under Messages with File Analysis Pending section, select **Quarantine** from the Action Applied to Message drop-down.  
The quarantined messages are stored in the File Analysis quarantine. See [Using File Analysis Quarantine, page 16-8](#).
- Step 4** (Optional) Under Messages with File Analysis Pending section, choose the following options:

- Whether to archive the original message. Archived messages are stored as an mbox-format log file in the `amparchive` directory on the appliance. The preconfigured AMP Archive (`amparchive`) log subscription is required.
- Whether to warn the end user by modifying the message subject, for example, “[WARNING: ATTACHMENT(S) MAY CONTAIN MALWARE].”
- Whether to add a custom header to provide granular controls to the administrator.

**Step 5** Submit and commit your changes.

---

#### Related Topics

- [Using File Analysis Quarantine, page 16-8](#)
- [Policy, Virus, and Outbreak Quarantines, page 30-1](#)

## Using File Analysis Quarantine

- [Edit File Analysis Quarantine Settings, page 16-8](#)
- [Manually Processing Messages in the File Analysis Quarantine, page 16-9](#)

### Edit File Analysis Quarantine Settings

#### Procedure

---

- Step 1** Select **Monitor > Policy, Virus, and Outbreak Quarantines**.
- Step 2** Click the **File Analysis** quarantine link.
- Step 3** Specify the retention period.
- Step 4** Specify the default action that AsyncOS must take after the retention period has passed.
- Step 5** If you do *not* want messages in this quarantine to be processed before the end of the Retention Period you specify, even when quarantine disk space is full, deselect **Free up space by applying default action on messages upon space overflow**.
- Step 6** If you select **Release** as the Default Action, optionally specify additional actions to apply to messages that are released before their retention period has passed:

| Option            | Information                                                                                                                                                                                                                                                                                                                                                              |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Modify Subject    | <p>Type the text to add and specify whether to add it to the beginning or the end of the original message subject.</p> <p>For example, you might want to warn the recipient that the message may contain malware attachments.</p> <p><b>Note</b> In order for a subject with non-ASCII characters to display correctly it must be represented according to RFC 2047.</p> |
| Add X-Header      | <p>An X-Header can provide a record of actions taken on a message. This can be helpful for example when handling inquiries about why a particular message was delivered.</p> <p>Enter a name and value.</p> <p>Example:</p> <p>Name =Inappropriate-release-early</p> <p>Value = True</p>                                                                                 |
| Strip Attachments | Stripping attachments protects against malware attachments in messages.                                                                                                                                                                                                                                                                                                  |

**Step 7** Specify the users who can access this quarantine:

| User                           | Information                                                                                                                                                                                                          |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local Users                    | <p>The list of local users includes only users with roles that can access quarantines.</p> <p>The list excludes users with Administrator privileges, because all Administrators have full access to quarantines.</p> |
| Externally Authenticated Users | You must have configured external authentication.                                                                                                                                                                    |
| Custom User Roles              | You see this option only if you have created at least one custom user role with quarantine access.                                                                                                                   |

**Step 8** Submit and commit your changes.

## Manually Processing Messages in the File Analysis Quarantine

### Procedure

**Step 1** Select **Monitor > Policy, Virus, and Outbreak Quarantines**.

**Step 2** In the row for File Analysis quarantine, click the blue number in the Messages column of the table.

**Step 3** Depending on your requirements, perform the following actions on messages:

- Delete
- Release
- Delay Scheduled Exit from quarantine
- Send a copy of messages to email addresses that you specify

## Centralized File Analysis Quarantine

For information about the centralized File Analysis quarantine, see [About Centralized Policy, Virus, and Outbreak Quarantines, page 30-10](#).

## X-Headers for File Reputation and Analysis

You can use X-Headers to mark messages with actions and results of message processing steps. You tag messages with X-Headers in mail policies, then use content filters to choose handling options and final actions for these messages.

Values are case-sensitive.

| Header Name            | Possible Values<br>(Case Sensitive) | Description                                                                                                                             |
|------------------------|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| X-Amp-Result           | Clean<br>Malicious<br>Unscannable   | Verdict applied to messages processed by the file reputation service.                                                                   |
| X-Amp-Original-Verdict | file unknown<br>verdict unknown     | Verdict before adjustment based on reputation threshold. This header exists only if the original verdict is one of the possible values. |
| X-Amp-File-Uploaded    | true<br>false                       | If any file attached to a message was sent for analysis, this header is "true."                                                         |

## Sending Notifications to End Users about Dropped Messages or Attachments

To send notifications to end users when a suspect attachment or its parent message has been dropped based on file reputation scanning, use an X-header or Custom Header and Content Filters.

## Advanced Malware Protection and Clusters

If you use centralized management, you can enable Advanced Malware Protection and mail policies at the cluster, group and machine level.

Feature keys must be added at the machine level.

## Ensuring That You Receive Alerts

Ensure that the appliance is configured to send you alerts related to Advanced Malware Protection.

You will receive alerts when:

| Alert Description                                                                                       | Type                              | Severity |
|---------------------------------------------------------------------------------------------------------|-----------------------------------|----------|
| Feature keys expire                                                                                     | (As is standard for all features) |          |
| The file reputation or file analysis service is unreachable                                             | System                            | Warning  |
| Communication with cloud services is established                                                        | System                            | Info     |
| The reputation and analysis engine is restarted by a watchdog service                                   | System                            | Info     |
| A file reputation verdict changes                                                                       | System                            | Info     |
| File types that can be sent for analysis have changed. You may want to enable upload of new file types. | System                            | Info     |
| Analysis of some file types is temporarily unavailable                                                  | System                            | Warning  |
| Analysis of all supported file types is restored after a temporary outage                               | System                            | Info     |

### Related Topics

- [Several Alerts About Failure to Connect to File Reputation or File Analysis Servers, page 16-14](#)
- [Taking Action When File Threat Verdicts Change, page 16-13](#)

## Configuring Centralized Reporting for Advanced Malware Protection Features

If you will centralize reporting on a Security Management appliance, see important configuration requirements in the Advanced Malware Protection sections in the email reporting chapter of the online help or user guide for your management appliance.

## File Reputation and File Analysis Reporting and Tracking

- [Identifying Files by SHA-256 Hash, page 16-11](#)
- [File Reputation and File Analysis Report Pages, page 16-12](#)
- [Viewing File Reputation Filtering Data in Other Reports, page 16-12](#)
- [About Message Tracking and Advanced Malware Protection Features, page 16-12](#)

## Identifying Files by SHA-256 Hash

Because filenames can easily be changed, the appliance generates an identifier for each file using a Secure Hash Algorithm (SHA-256). If an appliance processes the same file with different names, all instances are recognized as the same SHA-256. If multiple appliances process the same file, all instances of the file have the same SHA-256 identifier.



- In most reports, files are listed by their SHA-256 value (in an abbreviated format).

## File Reputation and File Analysis Report Pages

| Report                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Advanced Malware Protection | <p>Shows file-based threats that were identified by the file reputation service.</p> <p>For files with changed verdicts, see the AMP Verdict updates report. Those verdicts are not reflected in the Advanced Malware Protection report.</p> <p><b>Note</b> If one of the extracted files from a compressed or an archive file is malicious, only SHA value of the compressed or archive file is included in the Advanced Malware Protection report.</p>                                                                                                                                                                          |
| File Analysis               | <p>Displays the time and verdict (or interim verdict) for each file sent for analysis.</p> <p>To view more than 1000 File Analysis results, export the data as a .csv file.</p> <p>Drill down to view detailed analysis results, including the threat characteristics for each file.</p> <p>You can also search the cloud service for additional information about an SHA. The link is on the result details page.</p> <p><b>Note</b> If extracted files from a compressed or an archive file are sent for file analysis, only SHA values of these extracted files are included in the File Analysis report.</p>                  |
| AMP Verdict Updates         | <p>Lists the files processed by this appliance for which the verdict has changed since the message was received. For information about this situation, see <a href="#">File Threat Verdict Updates, page 16-1</a>.</p> <p>To view more than 1000 verdict updates, export the data as a .csv file.</p> <p>In the case of multiple verdict changes for a single SHA-256, this report shows only the latest verdict, not the verdict history.</p> <p>To view all affected messages for a particular SHA-256 within the maximum available time range (regardless of the time range selected for the report) click a SHA-256 link.</p> |

## Viewing File Reputation Filtering Data in Other Reports

Data for file reputation and analysis is available in other reports where relevant. A "Detected by Advanced Malware Protection" column may be hidden by default in applicable reports. To display additional columns, click the Columns link below the table.

## About Message Tracking and Advanced Malware Protection Features

When searching for file threat information in Message Tracking, keep the following points in mind:

- To search for malicious files found by the file reputation service, select **Advanced Malware Protection Positive** for the Message Event option in the Advanced section in Message Tracking.
- Message Tracking includes only information about file reputation processing and the original file reputation verdicts returned at the time a message was processed. For example, if a file was initially found to be clean, then a verdict update found the file to be malicious, only the clean verdict appears in Tracking results.

In Message Tracking details, the Processing Details section shows:

- The SHA-256 of each attachment in the message, and
- The final Advanced Malware Protection verdict for the message as a whole, and
- Any attachments which were found to contain malware.

No information is provided for clean or unscannable attachments.

- Verdict updates are available only in the AMP Verdict Updates report. The original message details in Message Tracking are not updated with verdict changes. To see messages that have a particular attachment, click a SHA-256 in the verdict updates report.
- Information about File Analysis, including analysis results and whether or not a file was sent for analysis, are available only in the File Analysis report.

Additional information about an analyzed file may be available from the cloud. To view any available File Analysis information for a file, select **Monitor > File Analysis** and enter the SHA-256 to search for the file. If the File Analysis service has analyzed the file from any source, you can see the details. Results are displayed only for files that have been analyzed.

If the appliance processed a subsequent instance of a file that was sent for analysis, those instances will appear in Message Tracking search results.

## Taking Action When File Threat Verdicts Change

### Procedure

- 
- |               |                                                                                                                                                                          |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | View the AMP Verdict Updates report.                                                                                                                                     |
| <b>Step 2</b> | Click the relevant SHA-256 link to view message tracking data for all messages that contained that file that may have been delivered to end users.                       |
| <b>Step 3</b> | Using the tracking data, identify the users that may have been compromised, as well as information such as the file names involved in the breach and sender of the file. |
| <b>Step 4</b> | Check the File Analysis report to see if this SHA-256 was sent for analysis, to understand the threat behavior of the file in more detail.                               |
- 

### Related Topics

- [File Threat Verdict Updates, page 16-1](#)

## Troubleshooting File Reputation and Analysis

- [Log Files, page 16-14](#)

- [Using Trace, page 16-14](#)
- [Several Alerts About Failure to Connect to File Reputation or File Analysis Servers, page 16-14](#)
- [Many Files Have Verdict "Unscannable", page 16-15](#)

## Log Files

In logs:

- `AMP` and `amp` refer to the file reputation service or engine.
- `Retrospective` refers to verdict updates.
- `VRT` and `sandboxing` refer to the file analysis service.

File reputation filtering and analysis events are logged in AMP Engine logs and Mail logs.

In the log message "Response received for file reputation query" possible values for "upload action" are:

- 0: The file is known to the reputation service; do not send for analysis.
- 1: Send
- 2: The file is known to the reputation service; do not send for analysis.

## Using Trace

Trace is not available for the file reputation filtering and analysis features. Instead, send a test message from an account outside your organization.

## Several Alerts About Failure to Connect to File Reputation or File Analysis Servers

**Problem** You receive several alerts about failures to connect to the file reputation or analysis services in the cloud. (A single alert may indicate only a transient issue.)

### Solution

- Ensure that you have met the requirements in [Requirements for Communication with File Reputation and Analysis Services, page 16-5](#).
- Check for network issues that may prevent the appliance from communicating with the cloud services.
- Increase the Query Timeout value:  
Select **Security Services > File Reputation and Analysis**. The Query Timeout value is in the Advanced settings area.

## Many Files Have Verdict "Unscannable"

**Problem** Many files have an "Unscannable" verdict.

**Solution** This can happen if there are connection issues between the appliance and the file reputation services in the cloud. Try the following:

- Check for issues on your network.
- Increase the timeout value on the **Security Services > File Reputation and Analysis** page.
- Check for an alert about communication errors between the cloud services and the appliance.





# Data Loss Prevention

---

- [Overview of Data Loss Prevention, page 17-1](#)
- [DLP Deployment Options, page 17-3](#)
- [System Requirements for Data Loss Prevention, page 17-4](#)
- [RSA Email DLP, page 17-4](#)
- [DLP Policies for RSA Email DLP, page 17-6](#)
- [RSA Enterprise Manager, page 17-23](#)
- [Message Actions, page 17-34](#)
- [Showing or Hiding Sensitive DLP Data in Message Tracking, page 17-38](#)
- [About Updating the DLP Engine and Content Matching Classifiers, page 17-39](#)
- [Working with DLP Incident Messages and Data, page 17-41](#)
- [Troubleshooting Data Loss Prevention, page 17-42](#)

## Overview of Data Loss Prevention

The Data Loss Prevention (DLP) feature secures your organization's proprietary information and intellectual property and enforces compliance with government regulations by preventing users from maliciously or unintentionally emailing sensitive data from your network. You define the types of data that your employees are not allowed to email by creating DLP policies that are used to scan outgoing messages for any data that may violate laws or corporate policies.

### Related Topics

- [Overview of the DLP Scanning Process, page 17-2](#)
- [How Data Loss Prevention Works, page 17-2](#)

## Overview of the DLP Scanning Process

|    | Action                                                                                                                                                                                                                                                                                                                              | More Information                                                                                                                                                                                                                             |
|----|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1. | A user in your organization sends an email message to a recipient outside of your organization.                                                                                                                                                                                                                                     | The Email Security appliance is a “gateway” appliance that processes messages that are entering or leaving your network.<br><br>Messages sent to other users within your network are not scanned.                                            |
| 2. | The Email Security appliance processes the message through the stages of its email “work queue” before it reaches the DLP scanning stage.                                                                                                                                                                                           | Pre-DLP-scanning processes ensure, for example, that the message includes no spam or malware.<br><br>To see where DLP processing occurs in the workqueue, see the workqueue flow diagram in <a href="#">Email Pipeline Flows, page 4-1</a> . |
| 3. | The appliance scans the message body, header, and attachments for sensitive content that you have identified in DLP Policies.                                                                                                                                                                                                       | See <a href="#">How Data Loss Prevention Works, page 17-2</a> .                                                                                                                                                                              |
| 4. | If sensitive content is found, the appliance takes action to protect the data, such as quarantining the message, dropping it, or delivering it with restrictions.<br><br>Otherwise, the message continues through the appliance’s work queue and if no issues are found, the Email Security appliance delivers it to the recipient. | You define the actions to be taken. See <a href="#">Message Actions, page 17-34</a> .                                                                                                                                                        |

## How Data Loss Prevention Works

When someone in your organization sends a message to a recipient outside your organization, the appliance determines which outgoing mail policy applies to the sender or recipient of that message, based on rules that you defined. The appliance evaluates the content of the message using the DLP policies that are specified in that outgoing mail policy.

Specifically, the appliance scans the message content (including headers and attachments) for text that matches words, phrases, predefined patterns such as social security numbers, or a regular expression that you identified as sensitive content in an applicable DLP policy.

The appliance also evaluates the context of disallowed content in order to minimize false positive matches. For example, a number matching a credit card number pattern is only a violation if it is accompanied by an expiration date, credit card company name (Visa, AMEX, etc.), or a person’s name and address.

If message content matches more than one DLP policy, the first matching DLP policy in the list applies, based on the order that you specified. If an outgoing mail policy has multiple DLP policies that use the same criteria to determine whether content is a violation, all policies use the result from a single content scan.

When potentially sensitive content appears in a message, the appliance assigns a risk factor score between 0 - 100 to the potential violation. This score indicates the likelihood that the message contains a DLP violation.

The appliance then assigns the severity level (such as Critical or Low) that you have defined for that risk factor score, and performs the message action that you have specified for that severity level in the applicable DLP Policy.

## DLP Deployment Options

| RSA Email DLP                                                                                                                                               | RSA Enterprise Manager                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| All DLP activities are handled by the Email Security appliance.                                                                                             | Third-party DLP management software from RSA that runs on a server and works with your Email Security appliance as “partnered devices.”<br><br><b>Note</b> RSA Enterprise Manager cannot be purchased from Cisco.                                                                                                                                         |
| Manages DLP policies on a single Email Security appliance, except in a cluster deployment.                                                                  | Manages DLP policies for multiple devices on the same network, including multiple Email Security appliances, from a centralized interface.                                                                                                                                                                                                                |
| You configure DLP policies on the Email Security appliance.                                                                                                 | Policies are configured in Enterprise Manager and pushed to the Email Security appliances on your network, for consistent DLP policies across your organization.                                                                                                                                                                                          |
| Includes over 100 DLP policy templates designed by RSA that your organization can use to define the sensitive data that your users cannot send using email. | Includes RSA’s DLP policy templates and integrates with RSA’s DLP Datacenter to use fingerprinting detection method for scanning source code and documents in certain DLP policies.<br><br>Fingerprinting is described in <a href="#">Fingerprinting, page 17-26</a> .                                                                                    |
| View and manage quarantined messages on the Email Security appliance or on a Security Management appliance.                                                 | Quarantined messages are stored on the Email Security appliance or on a Security Management appliance.<br><br>You can view quarantined messages on Enterprise Manager, on an Email Security appliance, or on a Security Management appliance.<br><br>You must manage quarantined messages (for example, delete or release them) using Enterprise Manager. |
| View and search reporting and tracking data on the Email Security appliance or on a Security Management appliance.                                          | View and search reporting and tracking data on Enterprise Manager, on the Email Security appliance, or on a Security Management appliance.                                                                                                                                                                                                                |
| —                                                                                                                                                           | Migrate the existing DLP configuration from your Email Security appliance to Enterprise Manager.                                                                                                                                                                                                                                                          |
| For more information, see <a href="#">RSA Email DLP, page 17-4</a> .                                                                                        | For more information, see <a href="#">RSA Enterprise Manager, page 17-23</a> .                                                                                                                                                                                                                                                                            |



**Note**

The following actions occur only on the Email Security appliance:

- Outgoing mail policy definition
- Message action definition
- DLP scanning

## System Requirements for Data Loss Prevention

Data Loss Prevention is supported on all supported C-Series and X-Series appliances except appliances using D-Mode licenses.

The RSA Enterprise Manager feature requires Enterprise Manager 9.0.

## RSA Email DLP

### Related Topics

- [How to Set Up Data Loss Prevention for Deployments Using RSA Email DLP, page 17-4](#)
- [Enabling Data Loss Prevention \(RSA Email DLP\), page 17-5](#)

## How to Set Up Data Loss Prevention for Deployments Using RSA Email DLP

Perform these steps in order:

|               | Do This                                                                                                                                                                                                                      | More Information                                                                                                                                                                                                                                                                                              |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Enable the DLP feature and choose RSA Email DLP as the deployment option.                                                                                                                                                    | <a href="#">Enabling Data Loss Prevention (RSA Email DLP), page 17-5</a>                                                                                                                                                                                                                                      |
| <b>Step 2</b> | Define the possible actions that can be taken for messages in which violations are found or suspected. For example, you can quarantine such messages.                                                                        | <a href="#">Message Actions, page 17-34</a>                                                                                                                                                                                                                                                                   |
| <b>Step 3</b> | Create DLP policies, which: <ul style="list-style-type: none"> <li>• identify the content that must not be emailed from your organization, and</li> <li>• specify which actions will be taken for each violation.</li> </ul> | Choose a method: <ul style="list-style-type: none"> <li>• <a href="#">Setting Up RSA Email DLP Using a Wizard, page 17-7</a></li> <li>• <a href="#">Creating a DLP Policy Using a Predefined Template, page 17-8</a></li> <li>• <a href="#">Creating a Custom DLP Policy (Advanced), page 17-9</a></li> </ul> |
| <b>Step 4</b> | Set the order of the DLP policies to determine which DLP policy is used to evaluate messages for DLP violations when the content could match more than one DLP policy.                                                       | <a href="#">Arranging the Order of the Email DLP Policies for Violation Matching, page 17-21</a>                                                                                                                                                                                                              |

|               | Do This                                                                                                                                         | More Information                                                                                                                                                                                                                               |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 5</b> | Ensure that you have created Outgoing Mail Policies for each group of senders and recipients whose messages will be scanned for DLP violations. | See <a href="#">Chapter 10, “Mail Policies.”</a><br>To further refine permitted and restricted message senders and recipients in individual DLP policies, see <a href="#">Filtering Messages for DLP Policies, page 17-20.</a>                 |
| <b>Step 6</b> | Specify which DLP policies apply to which senders and recipients by assigning DLP policies to Outgoing Mail Policies.                           | <a href="#">Associating DLP Policies with Outgoing Mail Policies, page 17-22</a>                                                                                                                                                               |
| <b>Step 7</b> | Configure settings for storage of and access to sensitive DLP information.                                                                      | <ul style="list-style-type: none"> <li>• <a href="#">Showing or Hiding Sensitive DLP Data in Message Tracking, page 17-38</a></li> <li>• <a href="#">Controlling Access to Sensitive Information in Message Tracking, page 32-5</a></li> </ul> |

## Enabling Data Loss Prevention (RSA Email DLP)

### Procedure

- Step 1** Select **Security Services > RSA Email DLP**.
- Step 2** Click **Enable**.
- Step 3** Scroll to the bottom of the license agreement page and click **Accept** to accept the agreement.



**Note** If you do not accept the license agreement, RSA Email DLP is not enabled on the appliance.

- Step 4** Under **Data Loss Prevention**, select **RSA Email DLP**.
- Step 5** Select the **Enable RSA Email Data Loss Prevention** check box.
- Step 6** (Recommended) For now, deselect the other options on this page.  
You can change these settings later, following instructions discussed elsewhere in this chapter.
- Step 7** Submit and commit your changes.

### What To Do Next

See [How to Set Up Data Loss Prevention for Deployments Using RSA Email DLP, page 17-4.](#)

### Related Topics

- [Showing or Hiding Sensitive DLP Data in Message Tracking, page 17-38](#)
- [Setting Up RSA Email DLP Using a Wizard, page 17-7](#)
- [About Updating the DLP Engine and Content Matching Classifiers, page 17-39](#)

# DLP Policies for RSA Email DLP

## Related Topics

- [DLP Policy Description, page 17-6](#)
- [Predefined DLP Policy Templates, page 17-6](#)
- [Setting Up RSA Email DLP Using a Wizard, page 17-7](#)
- [Creating a DLP Policy Using a Predefined Template, page 17-8](#)
- [Creating a Custom DLP Policy \(Advanced\), page 17-9](#)
- [About Defining Disallowed Content Using Content Matching Classifiers, page 17-10](#)
- [Filtering Messages for DLP Policies, page 17-20](#)
- [About Assessing Violation Severity, page 17-21](#)
- [Arranging the Order of the Email DLP Policies for Violation Matching, page 17-21](#)
- [Associating DLP Policies with Outgoing Mail Policies, page 17-22](#)
- [Important Information About Editing or Deleting DLP Policies, page 17-23](#)

## DLP Policy Description

A DLP policy includes:

- a set of conditions that determine whether an outgoing message contains sensitive data, and
- the actions to be taken when a message contains such data.

You specify how message content is evaluated, based on:

- Specific disallowed content or patterns of information. Depending on the policy, you may need to create a regular expression to search for identification numbers. See [About Defining Disallowed Content Using Content Matching Classifiers, page 17-10](#).
- A list of specific senders and recipients for filtering messages. See [Filtering by Senders and Recipients, page 17-20](#).
- A list of attachment file types for filtering messages. See [Filtering by Attachment Types, page 17-20](#).
- Settings that allow different actions to occur based on the severity of the violation. See [About Assessing Violation Severity, page 17-21](#).

You determine the message senders and recipients that each policy applies to when you enable DLP policies in Outgoing Mail Policies.

## Predefined DLP Policy Templates

To simplify creation of DLP policies, your appliance includes a large collection of predefined policy templates developed by RSA, Inc.

Template categories include:

- **Regulatory Compliance.** These templates identify messages and attachments that contain personally identifiable information, credit information, or other protected or non-public information.
- **Acceptable Use.** These templates identify messages sent to competitors or restricted recipients that contain sensitive information about an organization.
- **Privacy Protection.** These templates identify messages and attachments that contain identification numbers for financial accounts, tax records, or national IDs.
- **Intellectual Property Protection.** These templates identify popular publishing and design document file types that may contain intellectual property that an organization would want to protect.
- **Company Confidential.** These templates identify documents and messages that contain information about corporate accounting information and upcoming mergers and acquisitions.
- **Custom Policy.** This “template” lets you create your own policy from scratch using either content matching classifiers developed by RSA or violation identification criteria specified by your organization. This option is considered advanced and should be used only in the rare cases when the predefined policy templates do not meet the unique requirements of your network environment.

Some of these templates require customization.

## Setting Up RSA Email DLP Using a Wizard

The DLP Assessment Wizard helps you configure commonly-used DLP policies and enable them in the appliance’s default outgoing mail policy.



### Note

By default, DLP policies added using the DLP Assessment Wizard deliver all messages, regardless of the severity of detected DLP violations. You will need to edit the policies created using the wizard.

### Before You Begin

- Remove any existing DLP policies from the appliance. You can only use the DLP Assessment Wizard if there are no existing DLP policies on the appliance.
- If you need to detect messages that include student identification numbers or account numbers other than credit card numbers, US Social Security numbers, and US Drivers License numbers, create a regular expression that identifies those numbers. For more information, see [Regular Expressions for Identifying Identification Numbers](#), page 17-15.

### Procedure

- Step 1** Choose **Security Services > RSA Email DLP**.
- Step 2** Click **Edit Settings**.
- Step 3** Select the **Enable and configure DLP using the DLP Assessment Wizard** check box.
- Step 4** Click **Submit**.
- Step 5** Complete the wizard.  
Keep the following in mind:

- Any business that operates in California and owns or licenses computerized personally identifying information (PII) data for California residents, regardless of their physical location, is required to comply with **California SB-1386**. This law is one of the policy choices in the wizard.
- If you do not enter an email address to receive automatically-generated scheduled DLP Incident Summary report, the report will not be generated.
- When you review your configured settings, if you return to a step to make a change, you must proceed through the remaining steps until you reach the review page again. All settings that you previously entered will be remembered.
- When you complete the wizard, the Outgoing Mail Policies page displays, with your DLP policies enabled in the default outgoing mail policy. A summary of your DLP policy configuration is displayed at the top of the page.

**Step 6** Commit your changes.

---

#### What To Do Next

- (Optional) To edit these DLP policies, create additional policies, change the overall action on messages, or change the severity level settings, choose **Mail Policies > DLP Policy Manager**. For information, see [Creating a DLP Policy Using a Predefined Template](#), page 17-8, [Creating a Custom DLP Policy \(Advanced\)](#), page 17-9, and [Adjusting the Severity Scale](#), page 17-21.
- (Optional) To enable existing DLP policies for other outgoing mail policies, see [Using Outgoing Mail Policies to Assign DLP Policies to Senders and Recipients](#), page 17-23.

#### Related Topics

- [Creating a DLP Policy Using a Predefined Template](#), page 17-8
- [Creating a Custom DLP Policy \(Advanced\)](#), page 17-9

## Creating a DLP Policy Using a Predefined Template

#### Procedure

---

**Step 1** Select **Mail Policies > DLP Policy Manager**.

**Step 2** Click **Add DLP Policy**.

**Step 3** Click the name of a category to display a list of the available RSA Email DLP policy templates.



**Note** To view descriptions of each template, click **Display Policy Descriptions**.

---

**Step 4** Click **Add** for the RSA Email DLP policy template that you want to use.

**Step 5** (Optional) Change the predefined name and description of the template.

**Step 6** If the policy requires or recommends customizing one or more content matching classifiers, enter a regular expression to define the pattern of your organization's identification numbering system and a list of words or phrases related to the identification numbers that identify them as such or are typically associated with them.

For information, see:

- [About Defining Disallowed Content Using Content Matching Classifiers](#), page 17-10 and

- [Regular Expressions for Identifying Identification Numbers, page 17-15.](#)

**Note**

You cannot add or remove content matching classifiers for policies based on a predefined template.

**Step 7** (Optional) Apply the DLP policy only to messages with specific recipients, senders, attachment types, or previously-added message tags.

For more information, see [Filtering Messages for DLP Policies, page 17-20.](#)

You can separate multiple entries using a line break or a comma.

**Step 8** In the Severity Settings section:

- Choose an action to take for each level of violation severity.

For more information, see [About Assessing Violation Severity, page 17-21.](#)

- (Optional) Click **Edit Scale** to adjust the violation severity scale for the policy.

For more information, see [Adjusting the Severity Scale, page 17-21.](#)

**Step 9** Submit and commit your changes.

**Related Topics**

- [Setting Up RSA Email DLP Using a Wizard, page 17-7](#)
- [Creating a Custom DLP Policy \(Advanced\), page 17-9](#)

## Creating a Custom DLP Policy (Advanced)

**Note**

Creating custom policies is very complex; create custom policies only if the predefined DLP policy templates do not meet the needs of your organization.

You can create a custom DLP policy from scratch using the Custom Policy template and add either a predefined RSA content matching classifier or a custom classifier to the policy.

Custom policies can return a DLP violation if the content matches a single classifier or all classifiers, depending on how the policy is defined.

**Before You Begin**

Suggested: Define the criteria that identify a content violation. See [Creating a Content Matching Classifier for Custom DLP Policies, page 17-14.](#) You can also define these criteria from within this procedure.

**Procedure**

- Step 1** Select **Mail Policies > DLP Policy Manager.**
- Step 2** Click **Add DLP Policy.**
- Step 3** Click the name of the Custom Policy category.
- Step 4** Click **Add** for the Custom Policy template.

- Step 5** Enter a name and description for the policy.
- Step 6** Identify the content and context that constitute a DLP violation:
- Select a content matching classifier.
  - Click **Add**.
    - If you selected **Create a Classifier**, see [Creating a Content Matching Classifier for Custom DLP Policies, page 17-14](#).
    - Otherwise, the selected classifier is added to the table.
  - (Optional) Add additional classifiers to the policy.

For example, you might be able to eliminate known likely false positive matches by adding another classifier and selecting NOT.
  - If you added multiple classifiers: Choose an option in the table heading to specify whether **any** or **all** of the classifiers must match in order to count the instance as a violation.
- Step 7** (Optional) Apply the DLP policy only to messages with specific recipients, senders, attachment types, or previously-added message tags.
- For more information, see [Filtering Messages for DLP Policies, page 17-20](#).
- You can separate multiple entries using a line break or a comma.
- Step 8** In the Severity Settings section:
- Choose an action to take for each level of violation severity.

For more information, see [About Assessing Violation Severity, page 17-21](#).
  - (Optional) Click **Edit Scale** to adjust the violation severity scale for the policy.

For more information, see [Adjusting the Severity Scale, page 17-21](#).
- Step 9** Submit and commit your changes.
- 

#### Related Topics

- [Setting Up RSA Email DLP Using a Wizard, page 17-7](#)
- [Creating a DLP Policy Using a Predefined Template, page 17-8](#)
- [Duplicating a DLP Policy, page 17-11](#)

## About Defining Disallowed Content Using Content Matching Classifiers

Content matching classifiers define the content that cannot be emailed and optionally the context in which that content must occur in order to be considered a data loss prevention violation.

Suppose you want to prevent patient identification numbers from being emailed from your organization.

In order for the appliance to recognize these numbers, you must specify the patterns of the record numbering system used by your organization, using one or more regular expressions. You can also add a list of words and phrases that might accompany the record number as supporting information. If the classifier detects the number pattern in an outgoing message, it searches for the supporting information to verify that the pattern is an identification number and not a random number string. Including context matching information results in fewer false positive matches.

For this example, you might create a DLP policy that uses the HIPAA and HITECH template. This template includes the Patient Identification Numbers content matching classifier, which you can customize to detect a patient's identification number. To detect numbers in the pattern of 123-CL456789, you would enter the regular expression `[0-9]{3}\-[A-Z]{2}[0-9]{6}` for the classifier. Enter "Patient ID" for a related phrase. Finish creating the policy and enable it in an outgoing mail policy. Submit and commit your changes. Now, if the policy detects the number pattern in an outgoing message with the phrase "Patient ID" in close proximity to the number pattern, the DLP policy returns a DLP violation.

### About Using Content Matching Classifiers in DLP Policies

Many of the predefined DLP policy templates include content matching classifiers from RSA. Some of these classifiers require customization in order to identify the patterns that are used for data in your organization.

If you create a custom DLP policy, you can choose a predefined classifier or create one of your own.

### Related Topics

- [Content Matching Classifier Examples, page 17-12](#)
- [Creating a Content Matching Classifier for Custom DLP Policies, page 17-14](#)
- [Classifier Detection Rules for Identifying Sensitive Content \(Custom DLP Policies Only\), page 17-15](#)
- [Regular Expressions for Identifying Identification Numbers, page 17-15](#)
- [Using Custom Dictionaries of Sensitive DLP Terms \(Custom DLP Policies Only\), page 17-17](#)
- [Determiners of the Risk Factor of a Suspected Violation, page 17-18](#)
- [Viewing the Policies in Which Custom Content Classifiers are Used, page 17-20](#)



## Content Matching Classifier Examples

The following examples show how classifiers match message content:

- [Credit Card Number, page 17-12](#)
- [US Social Security Number, page 17-12](#)
- [ABA Routing Number, page 17-12](#)
- [US Drivers License, page 17-13](#)
- [US National Provider Identifier, page 17-13](#)
- [Student Records, page 17-13](#)
- [Corporate Financials, page 17-13](#)

### Credit Card Number

Several DLP policy templates include the Credit Card Number classifier. The credit card number itself is subject to various constraints, such as the pattern of digits and punctuation, the issuer-specific prefix, and the final check digit. The classifier requires additional supporting information to make a match, such as a second credit card number, an expiration date, or the name of the card issuer. This reduces the number of false positives.

Examples:

- 4999-9999-9999-9996 (No match because of no supporting information)
- 4999-9999-9999-9996 01/09 (Match)
- Visa 4999-9999-9999-9996 (Match)
- 4999-9999-9999-9996 4899 9999 9999 9997 (Match because of more than one credit card number)

### US Social Security Number

The US Social Security Number classifier requires a properly formatted number as well as supporting data, such as a date of birth, name, or the string `SSN`.

Examples:

- 321-02-3456 (No match because of no supporting information)
- 321-02-3456 July 4 (Match)
- 321-02-3456 7/4/1980 (Match)
- 321-02-3456 7/4 (No match)
- 321-02-3456 321-02-7654 (Match because of more than one SSN)
- SSN: 321-02-3456 (Match)
- Joe Smith 321-02-3456 (Match)
- 321-02-3456 CA 94066 (Match)

### ABA Routing Number

The ABA Routing Number classifier is similar to the Credit Card Number classifier.

Examples:

- 119999992 (No match because of no supporting information)

- routing 119999992 account 1234567 (Match)

## US Drivers License

Many policies use a US Drivers License classifier. By default, this classifier searches for drivers licenses for all 50 US states and the District of Columbia. Even US state-specific policies such as California AB-1298 and Montana HB-732 search for all 51 types of US drivers licenses. Thus, a predefined DLP policy template for a specific state, such as California SB 1386, uses the detection rules for all states and will return a DLP violation for data with a non-California driver license because it is still considered a privacy violation.

If you are concerned about false positives or appliance performance, you can limit searching to specific US states or no states by going to **Mail Policies > DLP Policy Manager** and clicking the **US Drivers Licenses** link in the **Advanced Settings** section.

The individual state classifiers match against the patterns for that state, and require the corresponding state name or abbreviation, and additional supporting data.

Examples:

- CA DL: C3452362 (Match because it has the correct pattern for the number and supporting data)
- California DL: C3452362 (Match)
- DL: C3452362 (No match because there is not enough supporting data)
- California C3452362 (No match because there is not enough supporting data)
- OR DL: C3452362 (No match because it is the incorrect pattern for Oregon)
- OR DL: 3452362 (Match because it is the correct pattern for Oregon)
- WV DL: D654321 (Match because it is the correct pattern for West Virginia)
- WV DL: G6543 (No match because it is the incorrect pattern for West Virginia)

## US National Provider Identifier

The US National Provider Identifier classifier scans for a US National Provider Identifier (NPI) numbers, which is a 10-digit number with a check digit.

Examples:

- NPI: 3459872347 (Match for NPI)
- 3459872347 (No match because of no supporting information)
- NPI: 3459872342 (No match because of incorrect check digit)

## Student Records

The predefined FERPA (Family Educational Rights and Privacy Act) DLP policy template uses the Student Records classifier. Combine it with a customized Student Identification Number classifier to detect specific student ID patterns for better accuracy.

Example:

- Joe Smith, Class Rank: 234, Major: Chemistry Transcript (Match)

## Corporate Financials

The predefined Sarbanes-Oxley (SOX) policy template uses the Corporate Financials classifier to search for non-public corporate financial information.

Examples:

2009 Cisco net sales, net income, depreciation (Match)

FORM 10-Q 2009 I.R.S. Employer Identification No. (Match)

## Creating a Content Matching Classifier for Custom DLP Policies

Custom classifiers that you create are added to the list of classifiers that you can use when creating custom DLP policies.

| Step          | Do This                                                                                                                                                                                                                                                                                             | Information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Understand how content matching classifiers are used to identify potential DLP violations.                                                                                                                                                                                                          | See: <ul style="list-style-type: none"> <li>• <a href="#">About Defining Disallowed Content Using Content Matching Classifiers</a>, page 17-10</li> <li>• <a href="#">Content Matching Classifier Examples</a>, page 17-12</li> </ul>                                                                                                                                                                                                                                                                                                   |
| <b>Step 2</b> | Select <b>Mail Policies &gt; DLP Policy Customizations</b> and click <b>Add Custom Classifier</b> .<br><br>Enter a classifier name and description.                                                                                                                                                 | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Step 3</b> | Enter a proximity and a minimum total score.                                                                                                                                                                                                                                                        | See <a href="#">Determiners of the Risk Factor of a Suspected Violation</a> , page 17-18                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Step 4</b> | Choose one of the following detection rule types and define the associated content matching criteria: <ul style="list-style-type: none"> <li>• words or phrases</li> <li>• text from a dictionary</li> <li>• a regular expression, or</li> <li>• an existing data loss prevention entity</li> </ul> | See: <ul style="list-style-type: none"> <li>• <a href="#">Classifier Detection Rules for Identifying Sensitive Content (Custom DLP Policies Only)</a>, page 17-15</li> <li>• <a href="#">Using Custom Dictionaries of Sensitive DLP Terms (Custom DLP Policies Only)</a>, page 17-17</li> <li>• <a href="#">Regular Expressions for Identifying Identification Numbers</a>, page 17-15</li> </ul> For information about Weight and Max Score, see <a href="#">Determiners of the Risk Factor of a Suspected Violation</a> , page 17-18. |
| <b>Step 5</b> | (Optional) Add additional rules by clicking <b>Add Rule</b> .                                                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Step 6</b> | If you include multiple rules, specify whether <b>All</b> or <b>Any</b> rules must match.                                                                                                                                                                                                           | This setting is at the top of the Rules section.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Step 7</b> | Submit and commit your changes.                                                                                                                                                                                                                                                                     | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

### What To Do Next

Use your custom content classifier in a custom DLP Policy. See [Creating a Custom DLP Policy \(Advanced\)](#), page 17-9.

**Related Topics**

- [Viewing the Policies in Which Custom Content Classifiers are Used, page 17-20](#)

**Classifier Detection Rules for Identifying Sensitive Content (Custom DLP Policies Only)**

Content matching classifiers require rules for detecting DLP violations in a message or document. Classifiers can use one or more of the following detection rules:

- **Words or Phrases.** A list of words and phrases for which the classifier should look. Separate multiple entries with a comma or line break.
- **Regular Expression.** A regular expression to define a search pattern for a message or attachment. You can also define a pattern to exclude from matching to prevent false positives. See [Regular Expressions for Identifying Identification Numbers, page 17-15](#) and [Examples of Regular Expressions for Identifying Identification Numbers, page 17-16](#) for more information.
- **Dictionary.** A dictionary of related words and phrases. Your appliance includes dictionaries created by RSA, or you can create your own. See [Using Custom Dictionaries of Sensitive DLP Terms \(Custom DLP Policies Only\), page 17-17](#).
- **Entity.** A predefined pattern that identifies common types of sensitive data, such as credit card numbers, addresses, social security numbers, or ABA routing numbers. For descriptions of the entities, go to **Mail Policies > DLP Policy Manager**, click **Add DLP Policy**, click **Privacy Protection**, then click **Display Policy Descriptions**.

**Regular Expressions for Identifying Identification Numbers**

Some policy templates require customization of one or more content matching classifiers, which involves creating a regular expression to search for identification numbers that may be linked to confidential information, such as a custom account number, patient identification number or Student ID. The style of regular expressions used for content matching classifiers is the **POSIX Basic Regular Expression** style regular expressions.

**Note**

Regular expressions are case sensitive, so they should include upper and lower case, such as `[a-zA-Z]`. If only certain letters are used, you can define the regular expression accordingly.

The less specific the pattern, such as an 8-digit number, the more likely you will want the policy to search for additional words and phrases to distinguish a random 8-digit number from an actual customer number.

Use the following table as a guide for creating regular expressions for classifiers:

| Element                                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Regular expression (abc)</b>         | <p>Regular expressions for classifiers match a string if the sequence of directives in the regular expression match any part of the string.</p> <p>For example, the regular expression <code>acc</code> matches the string <code>ACCOUNT</code> as well as <code>ACCT</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>[ ]</b>                              | <p>Use brackets to indicate a set of characters. Characters can be defined individually or within a range.</p> <p>For example, <code>[a-z]</code> matches all lowercase letters from <code>a</code> to <code>z</code>, while <code>[a-zA-Z]</code> matches all uppercase and lowercase letters from <code>A</code> to <code>Z</code>. <code>[xyz]</code> matches only the letters <code>x</code>, <code>y</code>, or <code>z</code>.</p>                                                                                                                                                                                                                                                 |
| <b>Backslash special characters (\)</b> | <p>The backslash character <i>escapes</i> special characters. Thus the sequence <code>\.</code> only matches a literal period, the sequence <code>\\$</code> only matches a literal dollar sign, and the sequence <code>\^</code> only matches a literal caret symbol.</p> <p>The backslash character also begins tokens, such as <code>\d</code>.</p> <p><b>Important Note:</b> The backslash is also a special escape character for the parser. As a result, if you want to include a backslash in your regular expression, you must use <i>two</i> backslashes — so that after parsing, only one “real” backslash remains, which is then passed to the regular expression system.</p> |
| <b>\d</b>                               | <p>Token that matches a digit (0-9). To match more than one digit, enter an integer in <code>{ }</code> to define the length of the number.</p> <p>For example, <code>\d</code> matches only a single digit such as <code>5</code>, but not <code>55</code>. Using <code>\d{2}</code> matches a number consisting of two digits, such as <code>55</code>, but not <code>5</code>.</p>                                                                                                                                                                                                                                                                                                    |
| <b>Number of repetitions {min,max}</b>  | <p>The regular expression notation that indicates the number of repetitions of the previous token is supported.</p> <p>For example, the expression <code>“\d{8}”</code> matches <code>12345678</code> and <code>11223344</code> but not <code>8</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Or ( )</b>                           | <p>Alternation, or the “or” operator. If <code>A</code> and <code>B</code> are regular expressions, the expression <code>“A B”</code> will match any string that matches either <code>“A”</code> or <code>“B”</code>. Can be used to combine number patterns in a regular expression.</p> <p>For example, the expression <code>“foo bar”</code> will match either <code>foo</code> or <code>bar</code>, but not <code>foobar</code>.</p>                                                                                                                                                                                                                                                 |

**Related Topics**

- [Examples of Regular Expressions for Identifying Identification Numbers, page 17-16](#)

**Examples of Regular Expressions for Identifying Identification Numbers**

Simple regular expressions that describe patterns of numbers and letters in identification or account numbers might look like the following:

- An 8-digit number: `\d{8}`
- Identification code with hyphens between sets of numbers: `\d{3}-\d{4}-\d{2}`
- Identification code that begins with a single letter that can be upper or lower case: `[a-zA-Z]\d{7}`
- Identification code that begins with three digits and is followed by nine uppercase letters:  
`\d{3}[A-Z]{9}`
- Using `|` to define two different number patterns to search for: `\d{3}[A-Z]{9}|\d{2}[A-Z]{9}-\d{2}`

## Using Custom Dictionaries of Sensitive DLP Terms (Custom DLP Policies Only)

AsyncOS comes with a set of predefined dictionaries from RSA Security Inc., but you can also create custom DLP dictionaries to specify terms for the DLP scanning feature to match.

You can create a custom DLP dictionary in several ways:

- [Adding Custom DLP Dictionaries Directly](#)
- [Creating DLP Dictionaries as Text Files](#) and then [Importing DLP Dictionaries](#).
- [Exporting DLP Dictionaries](#) from another Email Security appliance and then [Importing DLP Dictionaries](#).

### Adding Custom DLP Dictionaries Directly

#### Procedure

- 
- |               |                                                                                                                                                                                                                                         |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Select <b>Mail Policies &gt; DLP Policy Manager</b> .                                                                                                                                                                                   |
| <b>Step 2</b> | In the <b>Advanced Settings</b> section, click the link beside <b>Custom DLP Dictionaries</b> .                                                                                                                                         |
| <b>Step 3</b> | Click <b>Add Dictionary</b> .                                                                                                                                                                                                           |
| <b>Step 4</b> | Enter a name for the custom dictionary.                                                                                                                                                                                                 |
| <b>Step 5</b> | Enter new dictionary entries (words and phrases) into the list of terms.<br><br>Dictionary terms are case-sensitive and can contain non-ASCII characters.<br><br>When entering multiple entries, separate the entries with line breaks. |
| <b>Step 6</b> | Click <b>Add</b> .                                                                                                                                                                                                                      |
| <b>Step 7</b> | Submit and commit your changes.                                                                                                                                                                                                         |
- 

### Creating DLP Dictionaries as Text Files

You can create your own dictionary as a text file on your local machine and import it onto the appliance. Use line breaks for each term in the dictionary text file. Dictionary terms are case-sensitive and can contain non-ASCII characters.

### Exporting DLP Dictionaries



#### Note

Predefined DLP dictionaries cannot be exported.

### Procedure

- 
- Step 1** Select **Mail Policies > DLP Policy Manager**.
  - Step 2** Click the link for the **Custom DLP Dictionaries** section under Advanced Settings.
  - Step 3** Click **Export Dictionary**.
  - Step 4** Select a dictionary to export.
  - Step 5** Enter a file name for the dictionary.
  - Step 6** Choose where to save the exported dictionary, either on your local computer or in the configuration directory on the appliance.
  - Step 7** Select an encoding for the file.
  - Step 8** Click **Submit** and save the file.
- 

## Importing DLP Dictionaries

### Before You Begin

If you will import a file that you exported from a non-DLP dictionary on an Email Security appliance, you must first strip the weight values from the text file and convert any regular expressions to words or phrases.

### Procedure

- 
- Step 1** Select **Mail Policies > DLP Policy Manager**.
  - Step 2** In the **Advanced Settings** section, click the link beside **Custom DLP Dictionaries**.
  - Step 3** Click **Import Dictionary**.
  - Step 4** Select a file to import from either your local machine or the configuration directory on the appliance.
  - Step 5** Select an encoding.
  - Step 6** Click **Next**.  
A “Success” message appears and the imported dictionary is displayed in the Add Dictionary page. However, the process is not yet complete.
  - Step 7** Name and edit the dictionary.
  - Step 8** Click **Submit**.
- 

## Determiners of the Risk Factor of a Suspected Violation

When the appliance scans a message for DLP violations, it assigns a risk factor score to the message. This score indicates the likelihood that the message contains a DLP violation. A score of 0 means the message almost certainly does not contain a violation. A score of 100 means it almost certainly does contain a violation.

### For DLP Policies Based On Predefined Templates

You cannot view or modify risk factor scoring parameters for DLP policies created from predefined templates. However, if there are too many false positive matches for a particular DLP policy, you can adjust the severity scale for that policy. See [About Assessing Violation Severity, page 17-21](#). For policies based on templates that do not have a content matching classifier, such as the SOX (Sarbanes-Oxley) template, the scanning engine always returns a risk factor value of “75” when a message violates the policy.

### For Custom DLP Policies

When you create content matching classifiers for custom DLP policies, you specify values that are used to determine the risk factor score:

- **Proximity.** How close the rule matches must occur in the message or attachment to count as a violation. For example, if a numeric pattern similar to a social security number appears near the top of a long message and an address appears in the sender’s signature at the bottom, they are presumed to be unrelated and the data does not count as a match.
- **Minimum Total Score.** The minimum risk factor score required for sensitive content to be labeled a DLP violation. If the score of a message’s matches does not meet the minimum total score, its data is not considered sensitive.
- **Weight.** For each custom rule you create, you specify a “weight” to indicate the importance of the rule. A score is obtained by multiplying the number of detection rule matches by the weight of the rule. Two instances of a rule with a weight of 10 results in a score of 20. If one rule is more important for the classifier than the others, it should be assigned a greater weight.
- **Maximum Score.** A rule’s maximum score prevents a large number of matches for a low-weight rule from skewing the final score of the scan.

To calculate the risk factor, the classifier multiplies the number of matches for a detection rule by the weight of the rule. If this value exceeds the detection rule’s maximum score, the classifier uses the maximum score value. If the classifier has more than one detection rule, it adds the scores for all of its detection rules into a single value. The classifier maps the detection rules score (10 - 10000) on a scale of 10 -100 using the logarithmic scale shown in the following table to create the risk factor:

**Table 17-1**      *How Risk Factor Scores Are Calculated From Detection Rule Scores*

| Rule Scores | Risk Factor |
|-------------|-------------|
| 10          | 10          |
| 20          | 20          |
| 30          | 30          |
| 50          | 40          |
| 100         | 50          |
| 150         | 60          |
| 300         | 70          |
| 500         | 80          |
| 1000        | 90          |
| 10000       | 100         |



## Viewing the Policies in Which Custom Content Classifiers are Used

### Procedure

- 
- Step 1** Select **Mail Policies > DLP Policy Customizations**.
- Step 2** In the **Custom Classifiers** section, click the **Policies** link in the heading of the Custom Classifiers table.
- 

### Related Topics

- [Creating a Content Matching Classifier for Custom DLP Policies, page 17-14](#)

## Filtering Messages for DLP Policies

To improve performance or accuracy, you can limit a DLP policy to apply only to certain messages based on the following criteria:

| Option                              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Filtering by Senders and Recipients | <p>You can limit the DLP policy to apply to messages that do or do not include recipients or senders that you specify using one of the following:</p> <ul style="list-style-type: none"> <li>• Full email address: <code>user@example.com</code></li> <li>• Partial email address: <code>user@</code></li> <li>• All users in a domain: <code>@example.com</code></li> <li>• All users in a partial domain: <code>@.example.com</code></li> </ul> <p>Separate multiple entries using a line break or a comma.</p> <p>AsyncOS first matches the recipient or sender of an outgoing message to an outgoing mail policy, then matches the sender or recipient to the sender and recipient filters specified in the DLP policies enabled for that mail policy.</p> <p>For example, you might want to disallow all senders from sending a certain type of information, except to recipients in a partner domain. You would create a DLP policy for that information, including a filter that exempts all users in the partner domain, then include this DLP policy in an Outgoing Mail Policy that applies to all senders.</p> |
| Filtering by Attachment Types       | <p>You can limit the DLP policy to scanning only messages that do or do not include specific attachment types. Choose an attachment category, then a predefined file type, or specify file types that are not listed. If you specify a file type that is not predefined, AsyncOS searches for the file type based on the attachment's extension.</p> <p>You can also limit DLP scanning to attachments with a minimum file size.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Filtering by Message Tag            | <p>If you want to limit a DLP policy to messages containing a specific phrase, you can use a message or content filter to search outgoing messages for the phrase and insert a custom message tag into the message. For more information, see <a href="#">Content Filter Actions, page 11-9</a> and <a href="#">Chapter 9, "Using Message Filters to Enforce Email Policies."</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## About Assessing Violation Severity

When the DLP scanning engine detects a potential DLP violation, it calculates a risk factor score that represents the likelihood that the instance actually is a DLP violation. The policy compares the risk factor score to the Severity Scale defined in that policy in order to determine the severity level (for example, Low or Critical.) You specify the action to take for violations at each severity level (except Ignore, for which no action is ever taken.) You can adjust the risk factor scores required to reach each severity level.

### Related Topics

- [Adjusting the Severity Scale, page 17-21](#)

## Adjusting the Severity Scale

All policies have a default severity scale. You can adjust this scale for each policy.

For example, by default, a violation has a severity level of Critical if its risk factor score is between 90 and 100. However, for violations that match a particular policy, you may want increased sensitivity to potential data loss. For this DLP policy, you could change the Critical severity level to any violation with a risk factor score between 75 and 100.

### Procedure

- 
- |               |                                                                            |
|---------------|----------------------------------------------------------------------------|
| <b>Step 1</b> | Select <b>Mail Policies &gt; DLP Policy Manager</b> .                      |
| <b>Step 2</b> | Click the name of the policy to edit.                                      |
| <b>Step 3</b> | In the <b>Severity Settings</b> section, click <b>Edit Scale...</b>        |
| <b>Step 4</b> | Use the scale's arrows to adjust the scores for the severity levels.       |
| <b>Step 5</b> | Click <b>Done</b> .                                                        |
| <b>Step 6</b> | In the Severity Scale table, verify that your scores are as you want them. |
| <b>Step 7</b> | Click <b>Submit</b> .                                                      |
- 

### Related Topics

- [About Assessing Violation Severity, page 17-21](#)

## Arranging the Order of the Email DLP Policies for Violation Matching

If a DLP violation matches more than one of the DLP policies enabled in the outgoing mail policy, only the first matching DLP policy in the list is used.

### Procedure

- 
- |               |                                                                                            |
|---------------|--------------------------------------------------------------------------------------------|
| <b>Step 1</b> | On the DLP Policy Manager page, click <b>Edit Policy Order</b> .                           |
| <b>Step 2</b> | Click on the row for a policy you want to move and drag it to a new position in the order. |

**Step 3** Once you have finished reordering the policies, submit and commit your changes.

## Associating DLP Policies with Outgoing Mail Policies

### Related Topics

- [Associating DLP Policies with the Default Outgoing Mail Policy, page 17-22](#)
- [Using Outgoing Mail Policies to Assign DLP Policies to Senders and Recipients, page 17-23](#)

## Associating DLP Policies with the Default Outgoing Mail Policy

The default outgoing mail policy is used when no other outgoing mail policy matches the sender or a recipient.

**Figure 17-1** Default Outgoing Mail Policy with Enabled DLP Policies  
Outgoing Mail Policies

The screenshot shows the 'Find Policies' section with an 'Email Address' input field and radio buttons for 'Recipient' and 'Sender'. Below this is the 'Policies' section with an 'Add Policy...' button and a table of existing policies. The table has columns for Order, Policy Name, Anti-Spam, Anti-Virus, Content Filters, Outbreak Filters, DLP, and Delete. The 'Default Policy' row is highlighted in yellow and shows 'Not Available' for Anti-Spam and Anti-Virus, 'Disabled' for Content Filters, 'Not Available' for Outbreak Filters, and 'California SB-1386 Restricted Files' for DLP. At the bottom, there is a 'Key:' section with buttons for 'Default', 'Custom', and 'Disabled'.

| Order | Policy Name    | Anti-Spam     | Anti-Virus    | Content Filters | Outbreak Filters | DLP                                 | Delete |
|-------|----------------|---------------|---------------|-----------------|------------------|-------------------------------------|--------|
|       | Default Policy | Not Available | Not Available | Disabled        | Not Available    | California SB-1386 Restricted Files |        |

Key:

### Before You Begin

Complete all activities up to this point in the table in [How to Set Up Data Loss Prevention for Deployments Using RSA Email DLP, page 17-4](#). For example, ensure that you have created the DLP policies that you want to include in the default Outgoing Mail Policy.

### Procedure

- Step 1** Choose **Mail Policies > Outgoing Mail Policies**.
- Step 2** In the **Default Policy** row of the table, click the **Disabled** link in the **DLP** column.
- Step 3** Select **Enable DLP (Customize Settings)**.
- Step 4** Select the DLP policies to enable for the default outgoing mail policy.
- Step 5** Submit and commit your changes.

### What To Do Next

Choose the DLP policies for additional Outgoing Mail Policies. See [Using Outgoing Mail Policies to Assign DLP Policies to Senders and Recipients, page 17-23](#).

## Using Outgoing Mail Policies to Assign DLP Policies to Senders and Recipients

Specify which DLP policies apply to which senders and recipients by enabling them in outgoing mail policies. You can use DLP policies only in outgoing mail policies.

### Before You Begin

Configure the DLP policy settings for the default Outgoing Mail policy. See [Associating DLP Policies with the Default Outgoing Mail Policy, page 17-22](#).

### Procedure

- 
- |               |                                                                      |
|---------------|----------------------------------------------------------------------|
| <b>Step 1</b> | Choose <b>Mail Policies &gt; Outgoing Mail Policies</b> .            |
| <b>Step 2</b> | Click the link in the DLP column in any row of the table.            |
| <b>Step 3</b> | Select the DLP policies to associate with this outgoing mail policy. |
| <b>Step 4</b> | Submit your changes.                                                 |
| <b>Step 5</b> | Repeat as needed for other Outgoing Mail Policies.                   |
| <b>Step 6</b> | Commit your changes.                                                 |
- 

### What To Do Next

See [How to Set Up Data Loss Prevention for Deployments Using RSA Email DLP, page 17-4](#).

## Important Information About Editing or Deleting DLP Policies

| Action                | Information                                                                                                                                                                         |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Editing a DLP policy  | If you rename a policy, you must re-enable it in your outgoing mail policies.                                                                                                       |
| Deleting a DLP policy | If you delete a policy, you will receive a notification if the DLP policy is used in one or more outgoing mail policies. Deleting a DLP policy removes it from these mail policies. |

## RSA Enterprise Manager

### Related Topics

- [How Enterprise Manager and the Email Security Appliance Work Together, page 17-24](#)
- [Enterprise Manager Documentation, page 17-24](#)
- [How to Set up Data Loss Prevention in Deployments with RSA Enterprise Manager, page 17-24](#)
- [Migrating from RSA Email DLP to RSA Enterprise Manager, page 17-31](#)
- [Checking for DLP Policy Updates from Enterprise Manager, page 17-32](#)
- [RSA Enterprise Manager and Language Support, page 17-32](#)
- [Using Enterprise Manager with Clustered Appliances, page 17-32](#)

- [About Deleting and Disabling Policies in Enterprise Manager Deployments](#), page 17-33
- [Lost Connectivity Between the Email Security Appliance and Enterprise Manager](#), page 17-33
- [Switching from Enterprise Manager to RSA Email DLP](#), page 17-33

## How Enterprise Manager and the Email Security Appliance Work Together

When you enable RSA Enterprise Manager DLP on the Email Security appliance, the appliance sends the configuration to Enterprise Manager, which automatically adds the Email Security appliance as a partner device. The next time you open Enterprise Manager, the names and metadata of the outgoing mail policies and message actions that you configured on the Email Security appliance appear in Enterprise Manager, ready for you to use when configuring DLP policies. (Alternately, you can export existing DLP policies from the Email Security appliance to Enterprise Manager.)

After you configure DLP policies on Enterprise Manager, Enterprise Manager sends the DLP policies to the Email Security appliance. By default, all DLP policies pushed by Enterprise Manager are enabled on all devices they're pushed to, including Email Security appliances.

The Email Security appliance stores the DLP policies it receives from Enterprise Manager and uses them to scan outgoing messages for violations, and take action on any violations found. The Email Security appliance processes messages that are released for delivery, including encrypting the message if applicable. The Email Security appliance sends information about violations to Enterprise Manager for viewing and management.

### Related Topics

- [How Data Loss Prevention Works](#), page 17-2
- [DLP Deployment Options](#), page 17-3

## Enterprise Manager Documentation

For this deployment, you may need the following documentation from RSA Inc.:

- *Managing Partner Device DLP with Enterprise Manager* (technical note). Instructions on setting up Enterprise Manager and using it to manage the DLP features of partner devices, including Cisco Email Security appliances.
- *RSA DLP Network 9.0 Deployment Guide*. Instructions on deploying RSA DLP software on a network.
- *RSA DLP Network 9.0 User Guide*. Instructions for using the RSA DLP Network software, including how to use Enterprise Manager to manage partner DLP devices such as the Cisco Email Security appliance.

## How to Set up Data Loss Prevention in Deployments with RSA Enterprise Manager

Perform these steps in order:

|               | Do This                                                                                                                                                                                                                              | More Information                                                                                                                                                                                                                          |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Set up Enterprise Manager on your network and prepare for partnering with the Email Security appliance.                                                                                                                              | See RSA's documentation for DLP Datacenter, including the online help and the technical note <i>Managing Partner Device DLP with Enterprise Manager</i> .                                                                                 |
| <b>Step 2</b> | On the Email Security appliance, create Outgoing Mail Policies to determine which messages will be scanned for DLP violations.<br><br>Different policies can be assigned to different users or groups of users.                      | See <a href="#">Chapter 10, "Mail Policies."</a><br><br>Note:<br><br>The outgoing mail policy has an option to specify recipients. However, for deployments with Enterprise Manager, this information is not available from LDAP.         |
| <b>Step 3</b> | On the Email Security appliance, define the actions that can be taken for messages in which DLP violations are found or suspected.<br><br>For example, you can quarantine such messages.                                             | <a href="#">Message Actions, page 17-34</a>                                                                                                                                                                                               |
| <b>Step 4</b> | Obtain and upload certificates for secure communications between the Email Security appliance and Enterprise Manager.                                                                                                                | See <a href="#">(Recommended) Obtaining and Uploading Certificates for SSL Connections between Email Security Appliances and Enterprise Manager, page 17-26</a>                                                                           |
| <b>Step 5</b> | On the Email Security appliance, select RSA Enterprise Manager for the ESA's DLP Mode and configure the connection between the Email Security appliance and Enterprise Manager.                                                      | See <a href="#">Enabling Enterprise Manager DLP and Configuring the Connection with the Email Security Appliance, page 17-29.</a>                                                                                                         |
| <b>Step 6</b> | Provide the LDAP distinguished names of message senders to Enterprise Manager.                                                                                                                                                       | <a href="#">Using LDAP to Identify Message Senders for Enterprise Manager, page 17-30</a>                                                                                                                                                 |
| <b>Step 7</b> | If you will export DLP policies from the Email Security appliance and import them into Enterprise Manager, do so now.                                                                                                                | To export RSA Email DLP policies from the Email Security appliance, see <a href="#">Exporting DLP Policies from an Email Security Appliance, page 17-31.</a><br><br>To import the policies, see the RSA Enterprise Manager documentation. |
| <b>Step 8</b> | On Enterprise Manager, create DLP policies to: <ul style="list-style-type: none"> <li>• identify the types of content to be considered violations, and</li> <li>• specify which actions will be taken for each violation.</li> </ul> | Follow instructions for creating DLP policies in RSA's documentation for DLP Datacenter, including the online help and the technical note <i>Managing Partner Device DLP with Enterprise Manager</i> .                                    |
| <b>Step 9</b> | On Enterprise Manager, specify which DLP policies apply to which senders and recipients by associating DLP policies with Outgoing Mail Policies.                                                                                     | See <a href="#">About Associating Outgoing Mail Policies with DLP Policies in Enterprise Manager Deployments, page 17-31.</a>                                                                                                             |

|                | Do This                                                                                                                                                                               | More Information                                                                                                                                                                                                                               |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 10</b> | On Enterprise Manager, specify the order of the DLP policies.<br><br>When the appliance evaluates messages for DLP violations, it applies only the first matching policy in the list. | Order the DLP policies in Enterprise Manager. See the RSA Enterprise Manager documentation.                                                                                                                                                    |
| <b>Step 11</b> | On the Email Security appliance, configure settings for storage of and access to sensitive DLP information in Message Tracking.                                                       | <ul style="list-style-type: none"> <li>• <a href="#">Showing or Hiding Sensitive DLP Data in Message Tracking, page 17-38</a></li> <li>• <a href="#">Controlling Access to Sensitive Information in Message Tracking, page 32-5</a></li> </ul> |

**Related Topics**

- [Fingerprinting, page 17-26](#)
- [\(Recommended\) Obtaining and Uploading Certificates for SSL Connections between Email Security Appliances and Enterprise Manager, page 17-26](#)
- [Enabling Enterprise Manager DLP and Configuring the Connection with the Email Security Appliance, page 17-29](#)
- [Using LDAP to Identify Message Senders for Enterprise Manager, page 17-30](#)
- [About Associating Outgoing Mail Policies with DLP Policies in Enterprise Manager Deployments, page 17-31](#)

## Fingerprinting

If your Enterprise Manager deployment includes RSA's DLP Datacenter, you can enable fingerprinting. Fingerprinting improves detection of source code and sensitive documents including:

- Databases
- Full or partial text matches in the text of a document
- Full binary match, which is a bit-by-bit exact match of a file

If you enable fingerprinting, Enterprise Manager sends fingerprinting detection information to the Email Security appliance, and the Email Security appliance uses this information when scanning messages for Data Loss Prevention.

For more information about fingerprinting, see the Enterprise Manager documentation.

**Related Topics**

- [Enabling Enterprise Manager DLP and Configuring the Connection with the Email Security Appliance, page 17-29](#)

## (Recommended) Obtaining and Uploading Certificates for SSL Connections between Email Security Appliances and Enterprise Manager

If you want to use an SSL connection between the Email Security appliance and Enterprise Manager, you will need one or more certificates and signing keys from a recognized certificate authority to use for mutual authentication of the two machines.

When configuring the SSL connection, the Enterprise Manager server is the server and the Email Security appliance is the client.

Complete all of the following procedures:

- [Generating Client and Server Certificates using RSA's Certificate Tool, page 17-27](#)
- [Uploading a Certificate to the Email Security Appliance, page 17-28](#)
- [Uploading the Custom Certificate Authority File to the Email Security Appliance, page 17-28](#)
- [Generating a Certificate from the Email Security Appliance for Enterprise Manager, page 17-29](#)
- [Completing SSL Configuration, page 17-29](#)

## Generating Client and Server Certificates using RSA's Certificate Tool

RSA provides a certificate generation tool that you can use to generate a single .p12 file that you can use as both the server and client certificate for the connection. If you want to use different certificates for the appliance and the Enterprise Manager server, you must get them from another source.

This tool creates and stores two files on the Enterprise Manager server: the .p12 certificate file and a .pem certificate file. If you want to use the .p12 file, you must also import the .pem file onto the Email Security appliance as a certificate authority list.

For more information, see the RSA documentation.

### Procedure

**Step 1** Open a command prompt on the Enterprise Manager server.

**Step 2** Change to C:\Program Files\RSA\Enterprise Manager\etc.

**Step 3** Run the following command:

```
%JAVA_HOME%\bin/java" -cp ./emcerttool.jar
com.rsa.dlp.tem.X509CertGenerator -clientservercasigned -cacn <NAME OF CA PROVIDED DURING
INSTALL> -cakeystore catem-keystore -castorepass <PASSWORD FOR CA PROVIDED DURING
INSTALL> -cn <DEVICE_CN> -storepass <DEVICE STORE PASSWORD> -keystore <NAME OF DEVICE
STORE>
```



**Note** The common name of the certificate must be the hostname of the Email Security appliance.

If Enterprise Manager manages the connected Email Security appliances at the group or cluster level, each appliance requires a certificate with a Common Name that matches the hostname of that appliance.

A sample command may look like the following:

```
%JAVA_HOME%\bin/java" -cp ./emcerttool.jar
com.rsa.dlp.tem.X509CertGenerator -clientservercasigned -cacn emc-cisco
-cakeystore catem-keystore -castorepass esaem -cn ironport -storepass esaem
-keystore device-store
```



You can also use the following additional command-line switches:

```
-org <value in double quotes if it contains space>
-orgunit <value in double quotes if it contains space>
-title <value in double quotes if it contains space>
-validity <number of days>
```

This procedure outputs the <device-store>.p12 file to the same folder.

This .p12 file is the certificate that you will upload to the Email Security appliance.

You will also need:

- The .pem file from this folder, to import into the custom certificate authority list on the Email Security appliance.
- The Device Store password that you entered.

---

## Uploading a Certificate to the Email Security Appliance

### Before You Begin

Create a .p12 certificate file. You can use the procedure in [Generating Client and Server Certificates using RSA's Certificate Tool, page 17-27](#).

### Procedure

---

- |               |                                                                                                                                                                                                                                                                                                                                                              |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Select <b>Network &gt; Certificates</b> .                                                                                                                                                                                                                                                                                                                    |
| <b>Step 2</b> | Click <b>Add Certificate</b> .                                                                                                                                                                                                                                                                                                                               |
| <b>Step 3</b> | Select the <b>Import Certificate</b> option.                                                                                                                                                                                                                                                                                                                 |
| <b>Step 4</b> | Enter the path to the certificate file on your network or local machine.                                                                                                                                                                                                                                                                                     |
| <b>Step 5</b> | Enter the password for the file.                                                                                                                                                                                                                                                                                                                             |
| <b>Step 6</b> | Click <b>Next</b> to view the certificate's information.                                                                                                                                                                                                                                                                                                     |
| <b>Step 7</b> | Enter a name for the certificate. The Email Security appliance assigns the common name by default.<br>If Enterprise Manager manages the connected Email Security appliances at the group or cluster level, all certificates must have the same certificate name, even though the Common Name in each certificate is specific to each machine in the cluster. |
| <b>Step 8</b> | Submit and commit your changes.                                                                                                                                                                                                                                                                                                                              |
- 

## Uploading the Custom Certificate Authority File to the Email Security Appliance

### Before You Begin

Obtain the certificate authority file. If you generated a certificate using the procedure in [Generating Client and Server Certificates using RSA's Certificate Tool, page 17-27](#), this is the .pem file in the same folder as the .p12 certificate file.

### Procedure

- 
- Step 1** Select **Network > Certificates**.
- Step 2** In the Certificate Authorities section, click **Edit Settings**.
- Step 3** Click **Enable** for the Custom List.
- Step 4** Enter the full path to the custom list (the .pem file) on a local or network machine.
- Step 5** Submit and commit your changes.
- 

## Generating a Certificate from the Email Security Appliance for Enterprise Manager

If you prefer not to use the same certificate for client and server, you can generate a self-signed certificate from the Email Security appliance and upload it to Enterprise Manager. See [Creating a Self-Signed Certificate using the GUI, page 23-3](#).

## Completing SSL Configuration

You will complete the SSL configuration in “[Enabling Enterprise Manager DLP and Configuring the Connection with the Email Security Appliance, page 17-29](#).”

## Enabling Enterprise Manager DLP and Configuring the Connection with the Email Security Appliance

### Before You Begin

- Complete all steps prior to this step in the table in [How to Set up Data Loss Prevention in Deployments with RSA Enterprise Manager, page 17-24](#).
- If your deployment includes RSA’s DLP Datacenter, you can enable fingerprinting. For more information, see [Fingerprinting, page 17-26](#).

### Procedure

- 
- Step 1** Select **Security Services > RSA Email DLP** on the Email Security appliance.
- Step 2** If you have previously enabled Data Loss Prevention, click **Edit Settings** and then skip to [Step 5](#).
- Step 3** Click **Enable**.
- Step 4** Scroll to the bottom of the license agreement page and click **Accept** to accept the agreement.



**Note** If you do not accept the license agreement, Data Loss Prevention is not enabled on the appliance.

---

- Step 5** Under Data Loss Prevention, select **RSA Enterprise Manager**.
- Step 6** Enter the hostname for the Enterprise Manager server on your network that you want to use to manage DLP policies and 20000 for the port number. Separate the hostname and port number using a colon (:).
- Step 7** To use an SSL connection between the Email Security appliance and Enterprise Manager:
- Check the **Enable SSL Communication** check box
  - Select the Server Certificate. The server is Enterprise Manager.

- c. Select the Client Certificate. The client is the Email Security appliance.

You can use the same certificate for client and server.

**Step 8** (Optional) If your deployment includes RSA's DLP Datacenter, choose whether to enable fingerprinting to improve detection of source code, databases, and other documents.

**Step 9** (Optional) If message tracking is already enabled on your appliance, choose whether or not to enable matched content logging.

If you select this, the Email Security appliance logs DLP violations and AsyncOS displays the DLP violations and surrounding content in Message Tracking, including sensitive data such as credit card numbers and social security numbers.

**Step 10** Do not enable your appliance to automatically download updates to the DLP engine.

**Step 11** Submit and commit your changes.

The Email Security appliance sends the configuration to Enterprise Manager, which automatically adds the appliance as a partner device.

---

## Using LDAP to Identify Message Senders for Enterprise Manager

When the Email Security appliance sends DLP incident data to Enterprise Manager, the appliance must include the complete LDAP distinguished names in order to identify message senders. The appliance retrieves this information from an LDAP server.

### Before You Begin

- Complete all steps to this point in the table in [How to Set up Data Loss Prevention in Deployments with RSA Enterprise Manager, page 17-24](#). The User Distinguished Name Query option is not available unless you follow these instructions.
- Create an LDAP server profile on your Email Security appliance. See [Chapter 25, "LDAP Queries"](#) for more information.
- Create a query string that the appliance will use to retrieve the complete distinguished name unless you want to use the default query. For Active Directory servers, the default query string is `(proxyAddresses=smtp:{a})`. For OpenLDAP servers, the default query string is `(mail={a})`. You can define your own query and email attributes, including multiple attributes separated by commas.

### Procedure

---

**Step 1** Select **System Administration > LDAP** on the Email Security appliance.

**Step 2** Edit the profile for the LDAP server you want to use.

**Step 3** Select the check box for **User Distinguished Name Query**.

This option is available only if you have selected RSA Enterprise Manager as the DLP deployment option.

**Step 4** Enter a name for the query.

**Step 5** Enter the query string for retrieving the user distinguished name.

**Step 6** Click the button to test your query.

**Step 7** Submit and commit your changes.

---

## About Associating Outgoing Mail Policies with DLP Policies in Enterprise Manager Deployments

You will use Enterprise Manager to associate Outgoing Mail Policies with DLP policies, in order to specify which DLP policies apply to which senders and recipients. For information, see the RSA Enterprise Manager documentation. Enterprise Manager sends this information to the Email Security appliances for use when scanning messages.

Unlike with RSA Email DLP, outgoing mail policies cannot use the default mail policy's DLP settings when Enterprise Manager is enabled. If a mail policy is not specified for a DLP policy in Enterprise Manager, DLP scanning is not enabled on the mail policy.

## Migrating from RSA Email DLP to RSA Enterprise Manager

If you want to migrate your existing RSA Email DLP configuration to RSA Enterprise Manager, you can export your DLP configuration to a .zip file that you can upload to Enterprise Manager before switching the appliance from RSA Email DLP mode to RSA Enterprise Manager mode.

The Email Security appliance uses any existing local RSA Email DLP policies until it receives its first package of DLP policies from Enterprise Manager.

The Email Security appliance saves your existing RSA Email DLP policies in case you switch back to RSA Email DLP mode later.

### Related Topics

- [Exporting DLP Policies from an Email Security Appliance, page 17-31](#)
- [How to Set up Data Loss Prevention in Deployments with RSA Enterprise Manager, page 17-24](#)

## Exporting DLP Policies from an Email Security Appliance

You can export DLP policy configurations as a .zip file from an Email Security appliance and then import them to Enterprise Manager.

You can export DLP policies whether RSA Email DLP or RSA Enterprise Manager is selected as the DLP deployment mode.

### Procedure

---

**Step 1** Choose **Security Services > RSA Email DLP**.

**Step 2** Click **Export DLP Configuration**.

**Step 3** Enter a name for the .zip file and click **Export**.

Disabled DLP policies and DLP policies that are not assigned to an outgoing mail policy are not included.

**Note**

If the Email Security appliance is part of a cluster, the appliance only exports the policies from the lowest level of the cluster. For example, if there are DLP policies at both the cluster and machine level, the appliance only exports the DLP policies from the machine level.

**What To Do Next**

For information about importing the DLP policies into Enterprise Manager and distributing them to managed appliances, see the Enterprise Manager documentation.

## Checking for DLP Policy Updates from Enterprise Manager

Enterprise Manager periodically updates the DLP policies on the Email Security appliance.

To see the latest DLP policy update from Enterprise Manager, go to **Security Services > RSA Email DLP**.

**Related Topics**

- [Lost Connectivity Between the Email Security Appliance and Enterprise Manager, page 17-33](#)

## RSA Enterprise Manager and Language Support

The Email Security appliance displays any data it receives from RSA Enterprise Manager in the language that was used in Enterprise Manager. The appliance does not display this information in the language you selected for the appliance interface. This applies to DLP policies, content matching classifiers, dictionaries, and anything else created in Enterprise Manager that the appliance receives in the data package. For example, if the DLP policies and classifiers from Enterprise Manager were written in English but the interface of the Email Security appliance is displayed in French, the Email Security appliance displays the name and descriptions of the DLP policies and classifiers from Enterprise Manager in English. The rest of the interface remains in French.

## Using Enterprise Manager with Clustered Appliances

If you are using Enterprise Manager to manage the DLP policies for clustered Email Security appliances, be aware of the following:

- The Email Security appliance sends Enterprise Manager the outgoing mail policies and message actions from the lowest cluster level where these settings are configured. If these settings are configured differently at the cluster and machine level, the Email Security appliance sends Enterprise Manager the settings from the machine level. If you want to use the outgoing mail policies and message actions configured at a higher cluster level, delete the policies and actions defined at the lower levels that you do not want to use.
- The Email Security appliance uses the Data Loss Prevention mode specified at the lowest cluster level where this setting is configured. For example, if a clustered appliance is configured to use the local RSA Email DLP mode at machine level and RSA Enterprise Manager at the cluster level, the appliance uses RSA Email DLP for data loss prevention and does not communicate with Enterprise Manager.

## About Deleting and Disabling Policies in Enterprise Manager Deployments

### Deleting and Disabling DLP Policies

- To delete DLP policies, use Enterprise Manager.
- To disable or enable DLP policies, use the Email Security appliance. Go to **Mail Policies > DLP Policy Manager**.

Any outgoing mail policies associated with the disabled DLP policy will skip the policy when evaluating messages for DLP violations.

### Deleting Outgoing Mail Policies

If you try to delete an outgoing mail policy that is linked to a DLP policy, the Email Security appliance displays a message warning you that the mail policy is currently in use. If you delete the policy anyway, Enterprise Manager automatically unlinks the deleted outgoing mail policy from any DLP policy that used it. Other than not scanning for messages based on the configuration of the deleted mail policy, DLP scanning continues to work as before. The next DLP policy package sent to the Email Security appliance by Enterprise Manager will not include anything related to the deleted mail policy.

## Lost Connectivity Between the Email Security Appliance and Enterprise Manager

If connectivity between the Email Security appliance and Enterprise Manager is lost, any data that the appliance and Enterprise Manager cannot send is queued for delivery until the connection is restored. For the Email Security appliance, that means any data on messages containing possible DLP violations is queued. For Enterprise Manager, that means any data packages with new DLP policy information are queued. If the Email Security appliance does not receive updated DLP policy data from Enterprise Manager, the appliance continues to use the DLP policies it had previously received from Enterprise Manager.

### Related Topics

- [Enterprise Manager Disconnects the Email Security Appliance, page 17-42](#)

## Switching from Enterprise Manager to RSA Email DLP

If you want to go back to using RSA Email DLP for data loss prevention after using RSA Enterprise Manager, see [Enabling Data Loss Prevention \(RSA Email DLP\), page 17-5](#).

The Email Security appliance automatically reverts back to the RSA Email DLP policies it used before you configured it to use RSA Enterprise Manager mode. If the appliance did not use any local DLP policies when it was in RSA Email DLP mode, the appliance will continue to use the DLP policies from Enterprise Manager until you create a local DLP policy.

If you want to use local DLP policies similar to the ones on Enterprise Manager, you can recreate them using the DLP Policy Manager. The Email Security appliance does not automatically create new policies based on the ones used by Enterprise Manager and they cannot be imported from Enterprise Manager.

For information on creating DLP policies using the DLP Policy Manager, see [DLP Policies for RSA Email DLP, page 17-6](#).

For instructions on removing the Email Security appliance as a partner device in Enterprise Manager, see the RSA Enterprise Manager documentation.

# Message Actions

You specify primary and secondary actions that the Email Security appliance will take when it detects a possible DLP violation in an outgoing message. Different actions can be assigned for different violation types and severities.

Primary actions include:

- Deliver
- Drop
- Quarantine

Secondary actions include:

- Sending a copy to a policy quarantine if you choose to deliver the message. The copy is a perfect clone of the original, including the Message ID. Quarantining a copy allows you to test the RSA Email DLP system before deployment in addition to providing another way to monitor DLP violations. When you release the copy from the quarantine, the appliance delivers the copy to the recipient, who will have already received the original message.
- Encrypting messages. The appliance only encrypts the message body. It does not encrypt the message headers.
- Altering the subject header of messages containing a DLP violation.
- Adding disclaimer text to messages.
- Sending messages to an alternate destination mailhost.
- Sending copies (bcc) of messages to other recipients. (For example, you could copy messages with critical DLP violations to a compliance officer's mailbox for examination.)
- Sending a DLP violation notification message to the sender or other contacts, such as a manager or DLP compliance officer. See [Drafting DLP Notifications, page 17-36](#).

**Note**

These actions are not mutually exclusive: you can combine some of them within different DLP policies for various processing needs for different user groups. You can also configure different treatments based on the different severity levels in the same policy. For example, you may want to quarantine messages with critical DLP violations and send a notification to a compliance officer, but you may want to deliver messages with low severity levels.

**Related Topics**

- [Defining Actions to Take for DLP Violations \(Message Actions\), page 17-34](#)
- [Viewing and Editing Message Actions, page 17-36](#)
- [Drafting DLP Notifications, page 17-36](#)

## Defining Actions to Take for DLP Violations (Message Actions)

**Before You Begin**

- Create at least one dedicated quarantine to hold messages (or copies of messages) that violate DLP policies.

This can be a local quarantine on an Email Security appliance or a centralized quarantine on a Security Management appliance.

For deployments with Enterprise Manager:

- Set a timeout large enough for Enterprise Manager to complete its tasks.
- Consider automatic actions carefully; although quarantined messages must be managed in Enterprise Manager, the Email Security appliance still releases or deletes quarantined messages when the quarantine exceeds the allotted space.

For information, see [Chapter 30, “Policy, Virus, and Outbreak Quarantines.”](#)

- If you want to encrypt messages before delivery, make sure you have set up an encryption profile. See [Chapter 18, “Cisco Email Encryption.”](#)
- To include disclaimer text when delivering messages with DLP violations or suspected violations, specify disclaimer text in **Mail Policies > Text Resources**. For information, see [Disclaimer Template, page 21-12](#)
- To send a notification to the sender of a DLP violation or to another person such as a compliance officer, first create the DLP notification template. See [Drafting DLP Notifications, page 17-36](#).

### Procedure

- 
- Step 1** Select **Mail Policies > DLP Policy Customizations**.
- Step 2** In the **Message Actions** section, click **Add Message Action**.
- Step 3** Enter a name for the message action.
- Step 4** Enter a description of the message action.
- Step 5** Choose whether to drop, deliver, or quarantine messages containing DLP violations.



**Note** If you select Deliver, you can choose to have a copy of the message sent to a policy quarantine. The copy of the message is a perfect clone, including the Message ID.

---

- Step 6** If you want to encrypt the message upon delivery or its release from quarantine, select the **Enable Encryption** check box and select the following options:
- **Encryption Rule.** Always encrypts the message or only encrypt it if an attempt to send it over a TLS connection first fails.
  - **Encryption Profile.** Encrypts the message using the specified encryption profile and delivers it if you use a Cisco IronPort Encryption Appliance or a hosted key service.
  - **Encrypted Message Subject.** Subject for the encrypted message. Use the value is `$Subject` to keep the existing message subject.
- Step 7** If you select Quarantine as the action, choose the policy quarantine that you want to use for messages containing DLP violations.
- Step 8** Click **Advanced** if you want to modify the message using any of the following options:
- Add a custom header
  - Modify the message subject
  - Deliver it to an alternate host
  - Send a copy (bcc) to another recipient
  - Send a DLP notification message



**Step 9** Submit and commit your changes.

---

## Viewing and Editing Message Actions

### Procedure

---

**Step 1** Select **Mail Policies > DLP Policy Customizations**.

**Step 2** In the **Message Actions** section, choose an action:

| To                                                                                                                                                                                          | Do This                                                                                                                                                                       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| View the mail policies to which each action is assigned                                                                                                                                     | Click the <b>Policies</b> link in the heading of the Message Actions table.                                                                                                   |
| View the description that you entered for each action                                                                                                                                       | Click the <b>Description</b> link in the heading of the Message Actions table.                                                                                                |
| View or edit details of a Message Action                                                                                                                                                    | Click the name of the Message Action.                                                                                                                                         |
| Delete a Message Action                                                                                                                                                                     | Click the trash can icon next to the message action you want to delete.<br><br>A confirmation message notifies you if the message action is used in one or more DLP policies. |
| Duplicate a Message Action<br><br>You can use this feature to create a backup copy of a Message Action before changing it, or to use as a starting point for a new, similar Message Action. | Click the <b>Duplicate</b> icon next to the message action that you want to duplicate.                                                                                        |

**Step 3** Submit and commit any changes.

---

## Drafting DLP Notifications

Use this procedure to create a template for the notification that will be sent when an email message contains information that violates your organization's data loss prevention policies. You can send this notification to the sender of the message that violated DLP policy, or to another address, for example a manager or DLP Compliance officer.

### Before You Begin

- For deployments with RSA Enterprise Manager: You can configure either the Email Security appliance (Message Actions page) or Enterprise Manager (DLP policies) to send DLP violation notifications to users. To prevent duplicate notifications, set up notifications using one or the other, but not both.
- Familiarize yourself with the [DLP Notification Template Variable Definitions](#), page 17-37. You can use these variables to customize the notification with specific details about each violation.

### Procedure

- 
- Step 1** Select **Mail Policies > Text Resources**.
- Step 2** Click **Add Text Resource**.
- Step 3** For **Type**, select **DLP Notification Template**.  
DLP variables are not available for the plain Notification template.
- Step 4** Enter notification text and variables.  
The notification should inform its recipients that an outgoing message may contain sensitive data that violates your organization's data loss prevention policies.
- 

### What To Do Next

Specify this DLP notification template in a Message Action in a DLP policy in the DLP Policy Manager.

### Related Topics

- [DLP Notification Template Variable Definitions, page 17-37](#)

## DLP Notification Template Variable Definitions

Use the following variables to include specific information about each DLP violation in the notification.

| Variable               | Substituted With                                                                                                                                                     |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>\$DLPPolicy</b>     | Replaced by the name of the email DLP policy violated.                                                                                                               |
| <b>\$DLPSeverity</b>   | Replaced by the severity of violation. Can be "Low," "Medium," "High," or "Critical."                                                                                |
| <b>\$DLPRiskFactor</b> | Replaced by the risk factor of the message's sensitive material (score 0 - 100).                                                                                     |
| <b>\$To</b>            | Replaced by the message To: header (not the Envelope Recipient).                                                                                                     |
| <b>\$From</b>          | Replaced by the message From: header (not the Envelope Sender).                                                                                                      |
| <b>\$Subject</b>       | Replaced by the subject of the original message.                                                                                                                     |
| <b>\$Date</b>          | Replaced by the current date, using the format MM/DD/YYYY.                                                                                                           |
| <b>\$Time</b>          | Replaced by the current time, in the local time zone.                                                                                                                |
| <b>\$GMTimestamp</b>   | Replaced by the current time and date, as would be found in the Received: line of an email message, using GMT.                                                       |
| <b>\$MID</b>           | 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). |
| <b>\$Group</b>         | 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.      |
| <b>\$Reputation</b>    | Replaced by the SenderBase Reputation score of the sender. If there is no reputation score, it is replaced with "None".                                              |

| Variable                    | Substituted With                                                                                                                          |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>\$filenames</b>          | Replaced with a comma-separated list of the message's attachments' filenames.                                                             |
| <b>\$filetypes</b>          | Replaced with a comma-separated list of the message's attachments' file types.                                                            |
| <b>\$filesizes</b>          | Replaced with a comma-separated list of the message's attachment's file sizes.                                                            |
| <b>\$remotehost</b>         | Replaced by the hostname of the system that sent the message to the Cisco appliance.                                                      |
| <b>\$AllHeaders</b>         | Replaced by the message headers.                                                                                                          |
| <b>\$EnvelopeFrom</b>       | Replaced by the Envelope Sender (Envelope From, <MAIL FROM>) of the message.                                                              |
| <b>\$Hostname</b>           | Replaced by the hostname of the Cisco appliance.                                                                                          |
| <b>\$bodysize</b>           | Replaced by the size, in bytes, of the message.                                                                                           |
| <b>\$header['string']</b>   | Replaced by the value of the quoted header, if the original message contains a matching header. Note that double quotes may also be used. |
| <b>\$remoteip</b>           | Replaced by the IP address of the system that sent the message to the Cisco appliance.                                                    |
| <b>\$recvlistener</b>       | Replaced by the nickname of the listener that received the message.                                                                       |
| <b>\$dropped_filenames</b>  | Same as <code>\$filenames</code> , but displays list of dropped files.                                                                    |
| <b>\$dropped_filename</b>   | Returns only the most recently dropped filename.                                                                                          |
| <b>\$recvint</b>            | Replaced by the nickname of the interface that received the message.                                                                      |
| <b>\$timestamp</b>          | Replaced by the current time and date, as would be found in the Received: line of an email message, in the local time zone.               |
| <b>\$Time</b>               | Replaced by the current time, in the local time zone.                                                                                     |
| <b>\$orgid</b>              | Replaced by the SenderBase Organization ID (an integer value).                                                                            |
| <b>\$enveloperecipients</b> | Replaced by all Envelope Recipients (Envelope To, <RCPT TO>) of the message.                                                              |
| <b>\$dropped_filetypes</b>  | Same as <code>\$filetypes</code> , but displays list of dropped file types.                                                               |
| <b>\$dropped_filetype</b>   | Returns only the file type of the most recently dropped file.                                                                             |

## Showing or Hiding Sensitive DLP Data in Message Tracking

Both RSA Email DLP and RSA Enterprise Manager deployments offer the option to log the content that violates your DLP policies, along with the surrounding content, which can then be viewed in Message Tracking. This content may include sensitive data such as credit card numbers and social security numbers. You can opt not to log this content.

### Before You Begin

Enable Message Tracking. See [Chapter 29, “Tracking Messages.”](#)

**Procedure**

**Step 1** Select **Security Services > RSA Email DLP**.

**Step 2** Click **Edit Settings**.

| To                                             | Do This                                                       |
|------------------------------------------------|---------------------------------------------------------------|
| Include sensitive content in Message Tracking. | Select the <b>Enable Matched Content Logging</b> check box.   |
| Hide sensitive content from Message Tracking.  | Deselect the <b>Enable Matched Content Logging</b> check box. |

**Step 3** Submit and commit your changes.

**What To Do Next**

If you enable matched content logging, specify which administrative users can view this information. See [Controlling Access to Sensitive Information in Message Tracking, page 32-5](#).

## About Updating the DLP Engine and Content Matching Classifiers

Updates for the RSA DLP engine and the predefined content matching classifiers on your appliance are independent of updates for other security services.

**Related Topics**

- [Determining the Current Version of the RSA DLP Engine, page 17-39](#)
- [Caveats for DLP Updates, page 17-40](#)
- [Updating the DLP Engine and Content Matching Classifiers Manually, page 17-40](#)
- [Enabling Automatic Updates \(Not Recommended\), page 17-40](#)
- [DLP Updates on Centralized \(Clustered\) Appliances, page 17-41](#)
- [Rolling Back DLP Updates, page 17-41](#)

## Determining the Current Version of the RSA DLP Engine

**Procedure**

**Step 1** Select **Security Services > RSA Email DLP**.

**Step 2** Look in the **Current DLP Version Files** section.

## Caveats for DLP Updates

| Deployment Mode            | Caveat                                                                                                                                                                                                                                                                                                         |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| All                        | Cisco does not recommend enabling automatic updates. See <a href="#">Enabling Automatic Updates (Not Recommended)</a> , page 17-40                                                                                                                                                                             |
| RSA Email DLP              | DLP updates may change the content matching classifiers used by your existing local DLP policies. Cisco recommends that you manually download DLP updates to an appliance in a lab environment to test your DLP policies before updating appliances used in production.                                        |
| RSA Enterprise Manager DLP | Downloading DLP updates to your local appliance does not change the content matching classifiers used in your DLP policies, which are configured using Enterprise Manager. However, if you later switch your appliance to use RSA Email DLP, any existing local DLP policies will use the updated classifiers. |

## Updating the DLP Engine and Content Matching Classifiers Manually

### Before you Begin

See the following:

- [Caveats for DLP Updates](#), page 17-40
- (If applicable) [DLP Updates on Centralized \(Clustered\) Appliances](#), page 17-41

### Procedure

- 
- Step 1** Select **Security Services > RSA Email DLP**.
- Step 2** Click **Update Now** in the **Current DLP Version Files** section.
- This button is available only when there are new updates available for download.
- 

## Enabling Automatic Updates (Not Recommended)

Use this procedure to enable the appliance to check for and download updates at a regular interval.



### Note

Cisco recommends that you do not enable automatic updates. These updates may change the content matching classifiers used in your DLP policies. Instead, manually download DLP updates and test them in a lab environment before updating appliances used in production.

---

### Before You Begin

- On the **Security Settings > Service Updates** page, make sure you have enabled automatic updates and specified an update interval for all service updates.
- See [DLP Updates on Centralized \(Clustered\) Appliances](#), page 17-41.

### Procedure

- 
- |               |                                                       |
|---------------|-------------------------------------------------------|
| <b>Step 1</b> | Select <b>Security Services &gt; RSA Email DLP</b> .  |
| <b>Step 2</b> | Click <b>Edit Settings</b> .                          |
| <b>Step 3</b> | Select the <b>Enable automatic updates</b> check box. |
| <b>Step 4</b> | Submit and commit your changes.                       |
- 

## DLP Updates on Centralized (Clustered) Appliances

Note the following:

- You cannot enable automatic DLP updates for appliances in clustered deployments.
- DLP updates are performed at the level that DLP was configured. For example, if DLP was configured at cluster level, DLP updates must also be performed at that level.
- You can only roll back updates for appliances using the `dlprollback` CLI command at machine level.
- You can only check the status of an appliance's DLP engine using the `dlpstatus` CLI command at the machine level.

## Rolling Back DLP Updates

This procedure returns the system to using the previous DLP engine and content matching classifiers.



### Note

Rolling back DLP updates disables the DLP policies used in your mail policies.

### Before You Begin

See also [DLP Updates on Centralized \(Clustered\) Appliances, page 17-41](#).

### Procedure

- 
- |               |                                                        |
|---------------|--------------------------------------------------------|
| <b>Step 1</b> | In the CLI, use the <code>dlprollback</code> command.  |
| <b>Step 2</b> | Re-enable the DLP policies used in your mail policies. |
- 

## Working with DLP Incident Messages and Data



### Note

See also the documentation for Enterprise Manager and/or the Security Management appliance, as applicable to your deployment.

| To                                                                                                                                                             | Do This                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Search for messages containing DLP violations using criteria such as DLP policy name, violation severity, and action taken, and view details of messages found | See <a href="#">Chapter 29, “Tracking Messages.”</a><br>For Enterprise Manager deployments, you can also view messages in Enterprise Manager. See the Enterprise Manager documentation.                                                                                                                         |
| View or manage messages that have been quarantined as suspected DLP violations                                                                                 | See <a href="#">Working with Messages in Policy, Virus, or Outbreak Quarantines, page 30-10.</a><br>For Enterprise Manager deployments: You can view quarantined messages in Enterprise Manager or the Email Security appliance, but you must use Enterprise Manager to release or delete quarantined messages. |
| View a summary of DLP incidents                                                                                                                                | See information about DLP Incident Summary reports in <a href="#">Chapter 28, “Using Email Security Monitor.”</a>                                                                                                                                                                                               |
| View information about DLP violations discovered in outgoing mail                                                                                              | See information about DLP Incident reports in <a href="#">Chapter 28, “Using Email Security Monitor.”</a><br>For Enterprise Manager deployments, see the Enterprise Manager documentation.                                                                                                                      |
| Use Enterprise Manager to view DLP incident data and messages containing suspected violations                                                                  | See the documentation for Enterprise Manager.                                                                                                                                                                                                                                                                   |

**Related Topics**

- [Showing or Hiding Sensitive DLP Data in Message Tracking, page 17-38](#)
- [Controlling Access to Sensitive Information in Message Tracking, page 32-5](#)

## Troubleshooting Data Loss Prevention

- [Enterprise Manager Disconnects the Email Security Appliance, page 17-42](#)
- [RSA Email DLP Fails to Detect Violations in Email Attachments, page 17-43](#)

### Enterprise Manager Disconnects the Email Security Appliance

**Problem** , Enterprise Manager disconnects the Email Security appliance.

**Solution** The correct certificates have not been properly installed on the Email Security appliance. See [\(Recommended\) Obtaining and Uploading Certificates for SSL Connections between Email Security Appliances and Enterprise Manager, page 17-26.](#)

If you have a cluster or group deployment, check the Network > Certificates page on each appliance to verify that the certificate is the same on all.

**Related Topics**

- [Lost Connectivity Between the Email Security Appliance and Enterprise Manager, page 17-33](#)

## RSA Email DLP Fails to Detect Violations in Email Attachments

**Problem** When using predefined DLP policies, RSA Email DLP fails to detect violations in email attachments. This can be caused by the small value of the proximity parameter in the predefined DLP policies.



**Note** You cannot change the proximity of a predefined DLP policy.

**Solution** Do one of the following:

- Create a custom policy and adjust the proximity as required. See [Creating a Custom DLP Policy \(Advanced\)](#), page 17-9.
- Use RSA Enterprise Manager and the predefined policies in it. RSA Enterprise Manager allows you to fine-tune the predefined policy configurations such as proximity. See [RSA Enterprise Manager](#), page 17-23.
-







# Cisco Email Encryption

---

- [Overview of Cisco Email Encryption, page 18-1](#)
- [How to Encrypt Messages with a Local Key Server, page 18-2](#)
- [Encrypting Messages using the Email Security Appliance, page 18-4](#)
- [Determining Which Messages to Encrypt, page 18-8](#)
- [Inserting Encryption Headers into Messages, page 18-11](#)

## Overview of Cisco Email Encryption

AsyncOS supports using encryption to secure inbound and outbound email. To use this feature, you create an encryption profile that specifies characteristics of the encrypted message and connectivity information for the key server. The key server may either be:

- The Cisco Registered Envelope Service (managed service), or
- An Cisco Encryption appliance (locally managed server)

Next, you create content filters, message filters, and Data Loss Prevention policies to determine which messages to encrypt.

1. An outgoing message that meets the filter condition is placed in a queue on the Email Security appliance for encryption processing.
2. Once the message is encrypted, the key used to encrypt it is stored on the key server specified in the encryption profile and the encrypted message is queued for delivery.
3. If a temporary condition exists that prohibits the encryption of emails in the queue (i.e., temporary C-Series busyness or CRES unavailability), messages are re-queued and retried at a later time.



### Note

You can also set up the appliance to first attempt to send a message over a TLS connection before encrypting it. For more information, see [Using a TLS Connection as an Alternative to Encryption, page 18-9](#).

# How to Encrypt Messages with a Local Key Server

**Table 18-1**      *How to Encrypt Messages with a Local Key Server*


| Steps  | Do This                                                                                                                          | More Info                                                                                                                                                                                                                                                           |
|--------|----------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Set up the Cisco IronPort Encryption appliance on the network.                                                                   | See <a href="#">Chapter 3, “Setup and Installation.”</a>                                                                                                                                                                                                            |
| Step 2 | Enable message encryption.                                                                                                       | <a href="#">Enabling Message Encryption on the Email Security Appliance, page 18-4.</a>                                                                                                                                                                             |
| Step 3 | Specify the encryption key server to use and the security settings for the encrypted messages by creating an encryption profile. | <a href="#">Configuring How a Key Service Handles Encrypted Messages, page 18-4.</a>                                                                                                                                                                                |
| Step 4 | Define the conditions that messages must meet in order for the appliance to encrypt them.                                        | <a href="#">Determining Which Messages to Encrypt, page 18-8.</a>                                                                                                                                                                                                   |
| Step 5 | Determine when to encrypt messages in the email workflow.                                                                        | <ul style="list-style-type: none"> <li>• <a href="#">Encrypting and Immediately Delivering Messages using a Content Filter, page 18-9.</a></li> <li>or</li> <li>• <a href="#">Encrypting a Message upon Delivery using a Content Filter, page 18-10.</a></li> </ul> |
| Step 6 | (Optional) Flag messages for additional security.                                                                                | <a href="#">Inserting Encryption Headers into Messages, page 18-11.</a>                                                                                                                                                                                             |
| Step 7 | Define groups of users for whom you want to encrypt messages.                                                                    | Create a mail policy.<br>See <a href="#">Chapter 10, “Mail Policies.”</a>                                                                                                                                                                                           |
| Step 8 | Associate the encryption actions that you defined with the user groups you defined.                                              | Associate the content filter with the mail policy.<br>See <a href="#">Chapter 10, “Mail Policies.”</a>                                                                                                                                                              |

## Related Topics

- [Encryption Workflow, page 18-2](#)

## Encryption Workflow

When using email encryption, the Cisco Email Security appliance encrypts a message and stores the message key on a local key server or a hosted key service. When the recipient opens an encrypted message, the recipient is authenticated by the key service, and the decrypted message is displayed.

**Figure 18-1 Encryption Workflow**

The basic workflow for opening encrypted messages is:

1. When you configure an encryption profile, you specify the parameters for message encryption. For an encrypted message, the Email Security appliance creates and stores a message key on a local key server or on the hosted key service (Cisco Registered Envelope Service).
2. The recipient opens the secure envelope in a browser.
3. When a recipient opens an encrypted message in a browser, a password may be required to authenticate the recipient's identity. The key server returns the encryption key associated with the message.

**Note**

When opening an encrypted email message for the first time, the recipient is required to register with the key service to open the secure envelope. After registering, the recipient may be able to open encrypted messages without authenticating, depending on settings configured in the encryption profile. The encryption profile may specify that a password isn't required, but certain features will be unavailable.

4. The decrypted message is displayed.

# Encrypting Messages using the Email Security Appliance

To use encryption with the Email Security appliance, you must configure an encryption profile. You can enable and configure an encryption profile using the `encryptionconfig` CLI command, or via Security Services > Cisco IronPort Email Encryption in the GUI.



## Note

If PXE and S/MIME encryption is enabled on the appliance, AsyncOS encrypts messages using S/MIME first, and then using PXE.

## Related Topics

- [Enabling Message Encryption on the Email Security Appliance, page 18-4](#)
- [Configuring How a Key Service Handles Encrypted Messages, page 18-4](#)
- [Configuring the Default Locale of the Envelope, page 18-7](#)
- [Updating to the Latest Version of the PXE Engine, page 18-8](#)

## Enabling Message Encryption on the Email Security Appliance

### Procedure

**Step 1** Click **Security Services > Cisco IronPort Email Encryption**.

**Step 2** Click **Enable**.

**Step 3** (Optional) Click **Edit Settings** to configure the following options:

- The maximum message size to encrypt. Cisco's recommended message size is 10 MB. The maximum message size the appliance will encrypt is 25 MB.



## Note

Encrypting messages larger than the recommended 10 MB limit may slow down the performance of the appliance.

If you are using the Cisco Registered Envelope Service, message recipients will be unable to reply to an encrypted message that has attachments larger than 10 MB.

- Email address of the encryption account administrator. When you provision an Encryption Profile, this email address is registered automatically with the encryption server.
- Configure a proxy server.

## Configuring How a Key Service Handles Encrypted Messages

You can create one or more encryption profiles if you use a key service. You might want to create different encryption profiles if you want to use different levels of security for different groups of email. For example, you might want messages containing sensitive material to be sent with high security, but other messages to be sent with medium security. In this case, you might create a high security encryption profile to associate with the messages containing certain key words (such as 'confidential'), and create another encryption profile for other outgoing messages.

You can assign an encryption profile to a custom user role to allow delegated administrators assigned to that role to use the encryption profile with their DLP policies and content filters. Only administrators, operators, and delegated users can use encryption profiles when configuring DLP policies and content filters. Encryption profiles that are not assigned to a custom role are available for use by all delegated administrators with mail or DLP policy privileges. See [Distributing Administrative Tasks](#) for more information.

**Note**

You can configure multiple encryption profiles for a hosted key service. If your organization has multiple brands, this allows you to reference different logos stored on the key server for the PXE envelopes.

An encryption profile stores the following settings:

- **Key server settings.** Specify a key server and information for connecting to that key server.
- **Envelope settings.** Specify details about the message envelope, such as the level of security, whether to return read receipts, the length of time a message is queued for encryption before it times out, the type of encryption algorithm to use, and whether to enable a decryption applet to run on the browser.
- **Message settings.** Specify details about messages, such as whether to enable secure message forwarding and secure Reply All.
- **Notification settings.** Specify the notification template to use for text and HTML notifications, as well as encryption failure notifications. You create the templates in text resources and select the templates when creating the encryption profile. You can also localize envelopes and specify a message subject for encryption failure notifications. For more information about notifications, see [Encryption Notification Templates, page 21-24](#) and [Bounce and Encryption Failure Notification Templates, page 21-23](#).

**Procedure**

- 
- Step 1** In the Email Encryption Profiles section, click **Add Encryption Profile**.
- Step 2** Enter a name for the Encryption Profile.
- Step 3** Click the **Used By (Roles)** link, select the custom user role you want to have access to the encryption profile, and click **OK**.
- Delegated administrators assigned to this custom role can use the encryption profile for any DLP policies and content filters for which they are responsible.
- Step 4** In the Key Server Settings section, select from the following key servers:
- Cisco Encryption appliance (in network)
  - Cisco Registered Envelope Service (hosted key service)
- Step 5** If you select the Cisco Encryption appliance (local key service), enter the following settings:
- **Internal URL.** This URL is used by the Cisco Email Security appliance to contact the in-network Cisco Encryption appliance.
  - **External URL.** This URL is used when the recipient's message accesses keys and other services on the Cisco Encryption appliance. The recipient uses this URL to make inbound HTTP or HTTPS requests.
- Step 6** If you select the Cisco Registered Envelope Service, enter the URL for the hosted key service. The key service URL is `https://res.cisco.com`.

- Step 7** Click **Advanced** under Key Server Settings to specify whether to use HTTP or HTTPS for transferring the envelope's encrypted payload when the recipient opens the envelope. Choose from one of the following:
- **Use the Key Service with HTTP.** Transfers the encrypted payload from the key service using HTTP when the recipient opens the envelope. If you are using Cisco Registered Envelope Service, this is the URL you specified in [Step 6](#). If you are using the Cisco Encryption appliance, this is the external URL you specified in [Step 5](#).  
Since the payload is already encrypted, transporting it over HTTP is safe and faster than sending over HTTPS. This provides better performance than sending image requests over HTTPS.
  - **Use the Key Service with HTTPS.** Transfers the encrypted payload from the key service using HTTPS when the recipient opens the envelope. If you are using Cisco Registered Envelope Service, this is the URL you specified in [Step 6](#). If you are using the Cisco Encryption appliance, this is the external URL you specified in [Step 5](#).
  - **Specify a separate URL for payload transport.** If you don't want to use the key server for your encrypted payload, you can use another URL and specify whether to use HTTP or HTTPS for the payload transfer.
- Step 8** In the Envelope Settings section, select the level of message security:
- **High Security.** The recipient must always enter a password to open encrypted messages.
  - **Medium Security.** The recipient does not need to enter credentials to open the encrypted message if the recipient credentials are cached.
  - **No Password Required.** This is the lowest level of encrypted message security. The recipient does not need to enter a password to open the encrypted message. You can still enable the read receipts, Secure Reply All, and Secure Message Forwarding features for envelopes that are not password-protected.
- Step 9** To enable users to open your organization's URL by clicking its logo, you can add a link to the logo. Choose from the following options:
- **No link.** A live link is not added to the message envelope.
  - **Custom link URL.** Enter the URL to add a live link to the message envelope.
- Step 10** (Optional) Enable read receipts. If you enable this option, the sender receives a receipt when recipients open the secure envelope.
- Step 11** (Optional) Click **Advanced** under Envelope Settings to configure the following settings:
- Enter the length of time (in seconds) that a message can be in the encryption queue before timing out. Once a message times out, the appliance bounces the message and sends a notification to the sender.
  - Select an encryption algorithm:
    - **ARC4.** ARC4 is the most common choice, providing strong encryption with minimal decryption delays for message recipients.
    - **AES.** AES provides stronger encryption but also takes longer to decrypt, introducing delays for recipients. AES is typically used in government and banking applications.
  - Enable or disable the decryption applet. Enabling this option causes the message attachment to be opened in the browser environment. Disabling this option causes message attachments to be decrypted at the key server. If you disable this option, messages may take longer to open, but are not dependent on the browser environment.
- Step 12** In the Message Settings section, do the following:
- To enable secure reply all feature, check the **Enable Secure Reply All** check box.

- To enable secure message forwarding feature, check the **Enable Secure Message Forwarding** check box.

**Step 13** (Optional) If you have selected Cisco Registered Envelope Service and this service supports localization of envelopes, enable localization of envelopes. In Notification Settings section, check the **Use Localized Envelope** check box.



**Note** If you enable localization of envelopes, you cannot select encrypted message HTML or text notification.

If you want to set the default locale of the envelope, see [Configuring the Default Locale of the Envelope, page 18-7](#).

**Step 14** Select the HTML and text notification templates.



**Note** The key server uses an HTML or text notification based on the recipient's email application. You must configure notifications for both.

Do the following:

- a. Select an HTML notification template. Choose from HTML notifications you configured in text resources. If you did not configure a template, the system uses the default template.
- b. Select a text notification template. Choose from text notifications you configured in text resources. If you did not configure a template, the system uses the default template.



**Note** These options are unavailable if you use localized envelopes.

**Step 15** Enter a subject header for encryption failure notifications. The appliance sends a notification if the encryption process times out.

**Step 16** Select an encryption failure notification template for the message body. Choose from an encryption failure notification template you configured in text resources. If you did not configure a template, the system uses the default template.

**Step 17** Submit and commit your changes.

**Step 18** If you use Cisco Registered Envelope Service, you must take the additional step of provisioning your appliance. Provisioning the appliance registers the encryption profile with the hosted key service. To provision the appliance, click the **Provision** button for the encryption profile you want to register.

## Configuring the Default Locale of the Envelope

The default locale of the envelope is English. If you have selected Cisco Registered Envelope Service and this service supports localization of envelopes, you can change the locale of the envelope to any one of the following:

- English
- French
- German



- Japanese
- Portuguese
- Spanish

#### Before You Begin

- Create an encryption profile with Cisco Registered Envelope Service as Key Service Type and envelope localization enabled. See [Configuring How a Key Service Handles Encrypted Messages](#), page 18-4.
- Make sure that Cisco Registered Envelope Service supports localization of envelopes.

#### Procedure

- 
- Step 1** Click **Security Services > Cisco IronPort Email Encryption**.
- Step 2** Open an existing encryption profile.
- Step 3** In the **Notification Settings** section, choose the locale from the **Localized Envelopes** drop-down list.
- Step 4** Click **Submit**.
- Step 5** Click **Commit Changes**.
- 

## Updating to the Latest Version of the PXE Engine

The Cisco Email Encryption Settings page displays the current versions of the PXE engine and the Domain Mappings file used by your appliance. You can use the **Security Services > Service Updates** page (or the `updateconfig` command in the CLI) to configure the Email Security appliance to automatically update the PXE engine. For more information, see [Service Updates](#), page 33-17.

You can also manually update the engine using the **Update Now** button of the PXE Engine Updates section of IronPort Email Encryption Settings page (or the `encryptionupdate` command in the CLI).

## Determining Which Messages to Encrypt

After you create an encryption profile, you need to create an outgoing content filter that determines which email messages should be encrypted. The content filter scans outgoing email and determines if the message matches the conditions specified. Once the content filter determines a message matches the condition, the Cisco Email Security appliance encrypts the message and sends the generated key to the key server. It uses settings specified in the encryption profile to determine the key server to use and other encryption settings.

You can also encrypt messages after they are released after Data Loss Prevention scanning. For more information, see [Defining Actions to Take for DLP Violations \(Message Actions\)](#), page 17-34.

#### Related Topics

- [Using a TLS Connection as an Alternative to Encryption](#), page 18-9
- [Encrypting and Immediately Delivering Messages using a Content Filter](#), page 18-9

- [Encrypting a Message upon Delivery using a Content Filter, page 18-10](#)

## Using a TLS Connection as an Alternative to Encryption

Based on the destination controls specified for a domain, your Email Security appliance can securely relay a message over a TLS connection instead of encrypting it, if a TLS connection is available. The appliance decides whether to encrypt the message or send it over a TLS connection based on the TLS setting in the destination controls (Required, Preferred, or None) and the action defined in the encryption content filter.

When creating the content filter, you can specify whether to always encrypt a message or to attempt to send it over a TLS connection first, and if a TLS connection is unavailable, to encrypt the message. [Table 18-2](#) shows you how an Email Security appliance will send a message based on the TLS settings for a domain's destination controls, if the encryption control filter attempts to send the message over a TLS connection first.

**Table 18-2**      *TLS Support on ESA Appliances*

| Destination Controls TLS Setting | Action if TLS Connection Available | Action if TLS Connection Unavailable |
|----------------------------------|------------------------------------|--------------------------------------|
| None                             | Encrypt envelope and send          | Encrypt envelope and send            |
| TLS Preferred                    | Send over TLS                      | Encrypt envelope and send            |
| TLS Required                     | Send over TLS                      | Retry/bounce message                 |

For more information about enabling TLS on destination controls, see [Configuring the Gateway to Receive Email](#).

## Encrypting and Immediately Delivering Messages using a Content Filter

### Before You Begin

- To understand the concept of building conditions for content filters, see [Overview of Content Filters, page 11-1](#).
- (Optional) See [Inserting Encryption Headers into Messages, page 18-11](#).

### Procedure

- 
- Step 1** Go to **Mail Policies > Outgoing Content Filters**.
  - Step 2** In the Filters section, click **Add Filter**.
  - Step 3** In the Conditions section, click **Add Condition**.
  - Step 4** Add a condition to filter the messages that you want to encrypt. For example, to encrypt sensitive material, you might add a condition that identifies messages containing particular words or phrases, such as "Confidential," in the subject or body.
  - Step 5** Click **OK**.
  - Step 6** Optionally, click **Add Action** and select **Add Header** to insert an encryption header into the messages to specify an additional encryption setting.
  - Step 7** In the Actions section, click **Add Action**.

**Step 8** Select **Encrypt and Deliver Now (Final Action)** from the **Add Action** list.

**Step 9** Select whether to always encrypt messages that meet the condition or to only encrypt messages if the attempt to send it over a TLS connection fails.

**Step 10** Select the encryption profile to associate with the content filter.

The encryption profile specifies settings about the key server to use, levels of security, formatting of the message envelope, and other message settings. When you associate an encryption profile with the content filter, the content filter uses these stored settings to encrypt messages.

**Step 11** Enter a subject for the message.

**Step 12** Click **OK**.

The content filter in [Figure 18-2](#) shows a content filter that searches for ABA content in the message body. The action defined for the content filter specifies that the email is encrypted and delivered.

**Figure 18-2 Encryption Content Filter**

The screenshot shows the configuration interface for a content filter. It is divided into three main sections: Content Filter Settings, Conditions, and Actions.

**Content Filter Settings:**

- Name:** sensitive\_content
- Currently Used by Policies:** No policies currently use this rule.
- Description:** encrypt messages that contain sensitive material
- Order:** 2 (of 2)

**Conditions:**

| Order | Condition    | Rule                         | Delete   |
|-------|--------------|------------------------------|----------|
| 1     | Message Body | only-body-contains("aba", 1) | [Delete] |

**Actions:**

| Order | Action                             | Rule                                       | Delete   |
|-------|------------------------------------|--------------------------------------------|----------|
| 1     | Encrypt and Deliver (Final Action) | encrypt ("encrypt_sensitive", "\$Subject") | [Delete] |

At the bottom, there are **Cancel** and **Submit** buttons.

**Step 13** After you add the encrypt action, click **Submit**.

**Step 14** Commit your changes.

### What To Do Next

After you add the content filter, you need to add the filter to an outgoing mail policy. You may want to enable the content filter on the default policy, or you may choose to apply the filter to a specific mail policy, depending on your organization's needs. For information about working with mail policies, see [Overview of Mail Policies, page 10-1](#).

## Encrypting a Message upon Delivery using a Content Filter

Create a content filter to encrypt a message on delivery, which means that the message continues to the next stage of processing, and when all processing is complete, the message is encrypted and delivered.

### Before You Begin

- To understand the concept of building conditions for content filters, see [Overview of Content Filters, page 11-1](#).
- (Optional) See [Inserting Encryption Headers into Messages, page 18-11](#).

### Procedure

- 
- Step 1** Go to **Mail Policies > Outgoing Content Filters**.
- Step 2** In the Filters section, click **Add Filter**.
- Step 3** In the Conditions section, click **Add Condition**.
- Step 4** Add a condition to filter the messages that you want to encrypt. For example, to encrypt sensitive material, you might add a condition that identifies messages containing particular words or phrases, such as “Confidential,” in the subject or body.
- Step 5** Click **OK**.
- Step 6** Optionally, click **Add Action** and select **Add Header** to insert an encryption header into the messages to specify an additional encryption setting.
- Step 7** In the Actions section, click **Add Action**.
- Step 8** Select **Encrypt on Delivery** from the **Add Action** list.
- Step 9** Select whether to always encrypt messages that meet the condition or to only encrypt messages if the attempt to send it over a TLS connection fails.
- Step 10** Select the encryption profile to associate with the content filter.
- The encryption profile specifies settings about the key server to use, levels of security, formatting of the message envelope, and other message settings. When you associate an encryption profile with the content filter, the content filter uses these stored settings to encrypt messages.
- Step 11** Enter a subject for the message.
- Step 12** Click **OK**.
- Step 13** After you add the encrypt action, click **Submit**.
- Step 14** Commit your changes.
- 

### What To Do Next

After you add the content filter, you need to add the filter to an outgoing mail policy. You may want to enable the content filter on the default policy, or you may choose to apply the filter to a specific mail policy, depending on your organization’s needs. For information about working with mail policies, see [Overview of Mail Policies, page 10-1](#).

## Inserting Encryption Headers into Messages

AsyncOS enables you to add encryption settings to a message by inserting an SMTP header into a message using either a content filter or a message filter. The encryption header can override the encryption settings defined in the associated encryption profile, and it can apply specified encryption features to messages.



#### Note

The Cisco Ironport Encryption appliance must be set up to handle flagged messages.

**Procedure**

- Step 1** Go to **Mail Policies > Outgoing Content Filters** or **Incoming Content Filters**.
- Step 2** In the Filters section, click **Add Filter**.
- Step 3** In the Actions section, click **Add Action** and select **Add/Edit Header** to insert an encryption header into the messages to specify an additional encryption setting.

For example, if you want a Registered Envelope to expire in 24 hours after you send it, type X-PostX-ExpirationDate as the header name and +24:00:00 as the header value.

**Related Topics**

- [Encryption Headers, page 18-12](#)
- [Encryption Headers Examples, page 18-14](#)
- For more information about creating an encryption content filter, see [Encrypting and Immediately Delivering Messages using a Content Filter, page 18-9](#).
- For information about inserting a header using a message filter, see [Using Message Filters to Enforce Email Policies](#).

## Encryption Headers

Table 18-3 displays the encryption headers that you can add to messages.

**Table 18-3**      **Email Encryption Headers**

| MIME Header                 | Description                                                                                                                                                             | Value                                                                                                                                          |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| X-PostX-Reply-Enabled       | Indicates whether to enable secure reply for the message and displays the Reply button in the message bar. This header adds an encryption setting to the message.       | A Boolean for whether to display the Reply button. Set to <code>true</code> to display the button. The default value is <code>false</code> .   |
| X-PostX-Reply-All-Enabled   | Indicates whether to enable secure “reply all” for the message and displays the Reply All button in the message bar. This header overrides the default profile setting. | A Boolean for whether to display Reply All button. Set to <code>true</code> to display the button. The default value is <code>false</code> .   |
| X-PostX-Forward-Enabled     | Indicates whether to enable secure message forwarding and displays the Forward button in the message bar. This header overrides the default profile setting.            | A Boolean for whether to display the Forward button. Set to <code>true</code> to display the button. The default value is <code>false</code> . |
| X-PostX-Send-Return-Receipt | Indicates whether to enable read receipts. The sender receives a receipt when recipients open the Secure Envelope. This header overrides the default profile setting.   | A Boolean for whether to send a read receipt. Set to <code>true</code> to display the button. The default value is <code>false</code> .        |

**Table 18-3**      **Email Encryption Headers**

| MIME Header                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Value                                                                                                                                                                                             |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| X-PostX-ExpirationDate           | <p>Defines a Registered Envelope's expiration date before sending it. The key server restricts access to the Registered Envelope after the expiration date. The Registered Envelope displays a message indicating that the message has expired. This header adds an encryption setting to the message.</p> <p>If you use Cisco Registered Envelope Service, you can log in to the website at <a href="http://res.cisco.com">http://res.cisco.com</a> and use the message management features to set, adjust, or eliminate the expiration dates of messages after you send them.</p> | A string value containing relative date or time. Use the +HH:MM:SS format for relative hours, minutes, and seconds, and the +D format for relative days. By default, there is no expiration date. |
| X-PostX-ReadNotificationDate     | <p>Defines the Registered Envelope's "read by" date before sending it. The local key server generates a notification if the Registered Envelope has not been read by this date. Registered Envelopes with this header do not work with Cisco Registered Envelope Service, only a local key server. This header adds an encryption setting to the message.</p>                                                                                                                                                                                                                       | A string value containing relative date or time. Use the +HH:MM:SS format for relative hours, minutes, and seconds, and the +D format for relative days. By default, there is no expiration date. |
| X-PostX-Suppress-Applet-For-Open | <p>Indicates whether to disable the decryption applet. The decryption applet causes message attachments to be opened in the browser environment. Disabling the applet causes the message attachment to be decrypted at the key server. If you disable this option, messages may take longer to open, but they are not dependent on the browser environment. This header overrides the default profile setting.</p>                                                                                                                                                                  | A Boolean for whether to disable the decryption applet. Set to <code>true</code> to disable the applet. The default value is <code>false</code> .                                                 |

**Table 18-3**      **Email Encryption Headers**

| MIME Header                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Value                                                                                                                                                                         |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| X-PostX-Use-Script                     | Indicates whether to send JavaScript-free envelopes. A JavaScript-free envelope is a Registered Envelope that does not include the JavaScript that is used to open envelopes locally on the recipient's computer. The recipient must use either the Open Online method or the Open by Forwarding method to view the message. Use this header if a recipient domain's gateway strips JavaScript and makes the encrypted message unopenable. This header adds an encryption setting to the message. | A Boolean for whether the JavaScript applet should be included or not. Set to <code>false</code> to send a JavaScript-free envelope. The default value is <code>true</code> . |
| X-PostX-Remember-Envelope-Key-Checkbox | Indicates whether to allow envelope-specific key caching for offline opening of envelopes. With envelope key caching, the decryption key for a particular envelope is cached on the recipient's computer when the recipient enters the correct password and selects the "Remember the password for this envelope" check box. After that, the recipient does not need to enter a password again to reopen the envelope on the computer. This header adds an encryption setting to the message.     | A Boolean for whether to enable envelope key caching and display the "Remember the password for this envelope" check box. The default value is <code>false</code> .           |

## Encryption Headers Examples

This section provides examples of encryption headers.

### Related Topics

- [Enabling Envelope Key Caching for Offline Opening, page 18-14](#)
- [Enabling JavaScript-Free Envelopes, page 18-15](#)
- [Enabling Message Expiration, page 18-15](#)
- [Disabling the Decryption Applet, page 18-15](#)

## Enabling Envelope Key Caching for Offline Opening

To send a Registered Envelope with envelope key caching enabled, insert the following header into the message:

```
X-PostX-Remember-Envelope-Key-Checkbox: true
```

The "Remember the password for this envelope" check box is displayed on the Registered Envelope.

## Enabling JavaScript-Free Envelopes

To send a Registered Envelope that is JavaScript-free, insert the following header into the message:

```
X-PostX-Use-Script: false
```

When the recipient opens the securedoc.html attachment, the Registered Envelope is displayed with an Open Online link, and the Open button is disabled.

## Enabling Message Expiration

To configure a message so that it expires 24 hours after you send it, insert the following header into the message:

```
X-PostX-ExpirationDate: +24:00:00
```

The recipient can open and view the content of the encrypted message during the 24-hour period after you send it. After that, the Registered Envelope displays a message indicating that the envelope has expired.

## Disabling the Decryption Applet

To disable the decryption applet and have the message attachment decrypted at the key server, insert the following header into the message:

```
X-PostX-Suppress-Applet-For-Open: true
```

**Note**

The message may take longer to open when you disable the decryption applet, but it is not dependent on the browser environment.







## S/MIME Security Services

---

- [Overview of S/MIME Security Services, page 19-1](#)
- [Signing, Encrypting, or Signing and Encrypting Outgoing Messages using S/MIME, page 19-4](#)
- [Verifying, Decrypting, or Decrypting and Verifying Incoming Messages using S/MIME, page 19-14](#)
- [S/MIME Certificate Requirements, page 19-20](#)

### Overview of S/MIME Security Services

Secure/Multipurpose Internet Mail Extensions (S/MIME) is a standards-based method for sending and receiving secure, verified email messages. S/MIME uses public/private key pair to encrypt or sign messages. This way,

- If the message is encrypted, only the message recipient can open the encrypted message.
- If the message is signed, the message recipient can validate the identity of the sender's domain and can be assured that the message has not been altered while in transit.

For more information about S/MIME, review the following RFCs:

- RFC 5750: Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 - Certificate Handling
- RFC 5751: Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 - Message Specification
- RFC 3369: Cryptographic Message Syntax

### S/MIME Security Services in Email Security Appliance

Organizations may want to communicate securely using S/MIME without requiring that all end users possess their own certificates. For such organizations, Email Security appliance supports S/MIME security services (signing, encryption, verification, and decryption) at the gateway level using certificates that identify the organization rather than the individual user.

Email Security appliance provides the following S/MIME security services for Business-to-Business (B2B) and Business-to-Consumer (B2C) scenarios:

- Sign, encrypt, or sign and encrypt messages using S/MIME. See [Signing, Encrypting, or Signing and Encrypting Outgoing Messages using S/MIME, page 19-4](#).

- Verify, decrypt, or decrypt and verify messages using S/MIME. See [Verifying, Decrypting, or Decrypting and Verifying Incoming Messages using S/MIME](#), page 19-14.


#### Related Topics

- [Understanding How S/MIME Security Services Works](#), page 19-2

## Understanding How S/MIME Security Services Works

- [Scenario: Business-to-Business](#), page 19-2
- [Scenario: Business-to-Consumer](#), page 19-3

### Scenario: Business-to-Business



Organizations A and B want all the messages communicated between them to be signed and encrypted using S/MIME. Organization A has configured Email Security appliance to perform S/MIME security services at the gateway level. Organization B has configured a third-party application to perform S/MIME security services at the gateway level.



#### Note


The current example assumes that organization B is using a third-party application to perform S/MIME security services. In the real world, this can be any application or appliance (including Email Security appliance) that can perform S/MIME security services at the gateway level.

**Organization A sending a message to Organization B:**

1. Bob (Organization A) uses an email client to send an unsigned and unencrypted message to Dave (Organization B).
2. Email Security appliance in the Organization A signs and encrypts the messages and sends it to Organization B.
3. The third-party application at the gateway of Organization B decrypts and verifies the message.
4. Dave receives an unencrypted and unsigned message.

**Organization B sending a message to Organization A:**

1. Dave (Organization B) uses an email client to send an unsigned and unencrypted message to Bob (Organization A).
2. The third-party application at the gateway of Organization B signs and encrypts the message and sends it to Organization A.
3. Email Security appliance in the Organization A decrypts and verifies the message.
4. Bob receives an unencrypted and unsigned message.

**Scenario: Business-to-Consumer**

Organizations A and B want all the messages communicated between them to be signed and encrypted using S/MIME. Organization A has configured Email Security appliance to perform S/MIME security services at the gateway level. Organization B has configured the email clients of all the users to perform S/MIME security services.

**Organization A sending a message to Organization B:**

1. Alice (Organization A) uses an email client to send an unsigned and unencrypted message to Erin (Organization B).
2. Email Security appliance in the Organization A signs and encrypts the messages and sends it to Organization B.
3. The email client in the Organization B decrypts and verifies the message and displays it to Erin.

**Organization B sending a message to Organization A:**

1. Erin (Organization B) uses the email client to sign and encrypt a message and sends it to Alice (Organization A).
2. Email Security appliance in the Organization A decrypts and verifies the message.
3. Alice receives an unencrypted and unsigned message.

## Signing, Encrypting, or Signing and Encrypting Outgoing Messages using S/MIME

- [S/MIME Signing and Encryption Workflow in Email Security Appliance, page 19-4](#)
- [How to Sign, Encrypt, or Sign and Encrypt Outgoing Messages using S/MIME, page 19-6](#)
- [Setting Up Certificates for S/MIME Signing, page 19-6](#)
- [Setting Up Public Keys for S/MIME Encryption, page 19-9](#)
- [Managing S/MIME Sending Profiles, page 19-10](#)
- [Determining Which Messages to Sign, Encrypt, or Sign and Encrypt, page 19-13](#)
- [Signing, Encrypting, or Signing and Encrypting and Immediately Delivering Messages using a Content Filter, page 19-13](#)
- [Signing, Encrypting, or Signing and Encrypting a Message upon Delivery using a Content Filter, page 19-14](#)

**Note**

You can use Email Security appliance to sign, encrypt, and sign and encrypt outgoing and incoming messages.

## S/MIME Signing and Encryption Workflow in Email Security Appliance

- [S/MIME Signing Workflow, page 19-4](#)
- [S/MIME Encryption Workflow, page 19-5](#)

### S/MIME Signing Workflow

The following process describes how Email Security appliance performs S/MIME signing.

1. Apply a hash algorithm to the message to create a message digest.
2. Encrypt the message digest using private key of the appliance's S/MIME certificate.

3. Create a PKCS7 signature with the encrypted message digest and public key of the appliance's S/MIME certificate.
4. Sign the message by attaching the PKCS7 signature to the message.
5. Send the signed message to the recipient.

## S/MIME Encryption Workflow

The following process describes how Email Security appliance performs S/MIME encryption.

1. Create a pseudo-random session key.
2. Encrypt the message body using the session key.
3. Encrypt the session key using the public key of the recipient's (gateway or consumer) S/MIME certificate.
4. Attach the encrypted session key to the message.
5. Send the encrypted message to the recipient.

**Note**

If PXE and S/MIME encryption is enabled on the appliance, Email Security appliance encrypts messages using S/MIME first, and then using PXE.

## How to Sign, Encrypt, or Sign and Encrypt Outgoing Messages using S/MIME

| Steps  | Do This                                                                                                                                                                                                                                                                                                                                                                                                                    | More Info                                                                                                                                                                                                                                                                                                                       |
|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Understand the S/MIME certificate requirements.                                                                                                                                                                                                                                                                                                                                                                            | See <a href="#">S/MIME Certificate Requirements</a> , page 19-20.                                                                                                                                                                                                                                                               |
| Step 2 | Depending on your requirements, do one of the following: <ul style="list-style-type: none"> <li>For S/MIME signing, set up an S/MIME signing certificate.</li> <li>For S/MIME encryption, set up the public key of the recipient's S/MIME certificate.</li> <li>For S/MIME signing and encryption, set up an S/MIME signing certificate and the public key of the recipient's S/MIME certificate, respectively.</li> </ul> | See: <ul style="list-style-type: none"> <li><a href="#">Setting Up Certificates for S/MIME Signing</a>, page 19-6</li> <li><a href="#">Setting Up Public Keys for S/MIME Encryption</a>, page 19-9</li> </ul>                                                                                                                   |
| Step 3 | Create a profile for signing, encrypting, or signing and encrypting messages.                                                                                                                                                                                                                                                                                                                                              | See <a href="#">Create an S/MIME Sending Profile for Signing, Encrypting, or Signing and Encrypting Messages</a> , page 19-11.                                                                                                                                                                                                  |
| Step 4 | Define the conditions that messages must meet in order for the appliance to sign, encrypt, or sign and encrypt them.                                                                                                                                                                                                                                                                                                       | See <a href="#">Determining Which Messages to Sign, Encrypt, or Sign and Encrypt</a> , page 19-13.                                                                                                                                                                                                                              |
| Step 5 | Determine when in the email workflow to sign, encrypt, or sign and encrypt messages.                                                                                                                                                                                                                                                                                                                                       | See: <ul style="list-style-type: none"> <li><a href="#">Signing, Encrypting, or Signing and Encrypting and Immediately Delivering Messages using a Content Filter</a>, page 19-13</li> <li><a href="#">Signing, Encrypting, or Signing and Encrypting a Message upon Delivery using a Content Filter</a>, page 19-14</li> </ul> |
| Step 6 | Define groups of users for whom you want to sign or encrypt messages.                                                                                                                                                                                                                                                                                                                                                      | Create a mail policy.<br>See <a href="#">Chapter 10, "Mail Policies."</a>                                                                                                                                                                                                                                                       |
| Step 7 | Associate the signing or encryption actions that you defined with the user groups you defined.                                                                                                                                                                                                                                                                                                                             | Associate the content filter with the mail policy.<br>See <a href="#">Chapter 10, "Mail Policies."</a>                                                                                                                                                                                                                          |



### Note

If you want to perform S/MIME signing, encryption, or signing and encryption using CLI, use the `smimeconfig` command. See *Cisco AsyncOS for Email CLI Reference Guide*.

## Setting Up Certificates for S/MIME Signing

You must set up an S/MIME certificate for signing messages. The Email Security appliance allows you to set up S/MIME signing certificates using one of the following methods:

- Create a self-signed S/MIME certificate using the appliance. See [Creating a Self-Signed S/MIME Certificate](#), page 19-7.

- Import an existing S/MIME certificate to the appliance. See [Importing an S/MIME Signing Certificate, page 19-8](#).

**Note**

Cisco recommends that you use self-signed S/MIME certificates for sending signed messages to the users within your organization or in a testing environment. For sending signed messages to external users or in a production environment, use a valid S/MIME certificate obtained from a trusted CA.

For understanding the certificate requirements for S/MIME, see [S/MIME Certificate Requirements, page 19-20](#).

## Creating a Self-Signed S/MIME Certificate

You can generate self-signed S/MIME certificates that are compliant to RFC 5750 (Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 - Certificate Handling) using the web interface or CLI.

**Note**

Cisco recommends that you use self-signed S/MIME certificates for sending signed messages to the users within your organization or in a testing environment.

### Procedure

- Step 1** Click **Network > Certificates**.
- Step 2** Click **Add Certificate**.
- Step 3** Choose **Create Self-Signed S/MIME Certificate**.
- Step 4** Enter the following information for the self-signed certificate:

|                                          |                                                                                                                                                                                                                                                                                                   |
|------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Common Name</b>                       | The fully qualified domain name.                                                                                                                                                                                                                                                                  |
| <b>Organization</b>                      | The exact legal name of the organization.                                                                                                                                                                                                                                                         |
| <b>Organizational Unit</b>               | Section of the organization.                                                                                                                                                                                                                                                                      |
| <b>City (Locality)</b>                   | The city where the organization is legally located.                                                                                                                                                                                                                                               |
| <b>State (Province)</b>                  | The state, county, or region where the organization is legally located.                                                                                                                                                                                                                           |
| <b>Country</b>                           | The two letter ISO abbreviation of the country where the organization is legally located.                                                                                                                                                                                                         |
| <b>Duration before expiration</b>        | The number of days before the certificate expires.                                                                                                                                                                                                                                                |
| <b>Subject Alternative Name(Domains)</b> | <p>If you configure this field, any user from the specified domain can send signed messages.</p> <p>Name of the domain from which you plan to send signed messages. Examples include <code>domain.com</code> and <code>*.domain.net</code>. For multiple entries, use a comma-separated list.</p> |
| <b>Subject Alternative Name(Email)</b>   | <p>If you configure this field, only the specified users can send signed messages.</p> <p>Email address of the user who is planning to send signed messages, for example, <code>user@somedomain.com</code>. For multiple entries, use a comma-separated list.</p>                                 |
| <b>Private Key Size</b>                  | Size of the private key to generate the certificate signing request (CSR).                                                                                                                                                                                                                        |



**Note**

An S/MIME signing certificate can contain both Subject Alternative Name (Domains) and Subject Alternative Name (Email).

- Step 5** Click **Next** to view the certificate and signature information.
- Step 6** Depending on your requirements, do the following:
- Enter a name for the certificate.
  - If you want to submit a CSR for the self-signed certificate to a certificate authority, click **Download Certificate Signing Request** to save the CSR in PEM format to a local or network machine.
- Step 7** Submit and commit your changes.

**Note**

Use the `certconfig` command to generate self-signed S/MIME certificates using CLI.

## Importing an S/MIME Signing Certificate

If you already have an S/MIME certificate for signing messages, you can add it to the appliance by importing it.

### Before You Begin

Make sure that the S/MIME certificate that you plan to import meets the requirements described in [S/MIME Certificate Requirements, page 19-20](#).

### Procedure

- Step 1** Click **Network > Certificates**.
- Step 2** Click **Add Certificate**.
- Step 3** Choose **Import Certificate**.
- Step 4** Enter the path to the certificate file on your network or local machine.
- Step 5** Enter the password for the file.
- Step 6** Click **Next** to view the certificate's information.
- Step 7** Enter a name for the certificate.
- Step 8** Submit and commit your changes.

**Note**

Use the `certconfig` command to import S/MIME certificates using CLI.

## Setting Up Public Keys for S/MIME Encryption

You must add the public key of the recipient's S/MIME certificate to the appliance for encrypting messages. Depending on your organizational policies and processes, you can use one of the following methods to add the public key to the appliance:

- Request the recipient to send the public key using an electronic channel, for example, email. You can then add the public key using the web interface or CLI.

For instructions to add the public key, see [Adding a Public Key for S/MIME Encryption, page 19-9](#).

- Enable public key harvesting using the web interface or CLI and request the recipient to send a signed message. The Email Security appliance can harvest the public key from the signed message.

For instructions to harvest public key from an incoming signed message, see [Harvesting Public Keys, page 19-9](#).

## Adding a Public Key for S/MIME Encryption

### Before You Begin

- Make sure that the public key meets the requirements described in [S/MIME Certificate Requirements, page 19-20](#).
- Make sure that the public key is in PEM format.

### Procedure

- 
- |               |                                               |
|---------------|-----------------------------------------------|
| <b>Step 1</b> | Click <b>Mail Policies &gt; Public Keys</b> . |
| <b>Step 2</b> | Click <b>Add Public Key</b> .                 |
| <b>Step 3</b> | Enter the name of the public key.             |
| <b>Step 4</b> | Enter the public key.                         |
| <b>Step 5</b> | Submit and commit your changes.               |
- 



#### Note

Use the `smimeconfig` command to add public keys using CLI.

## Harvesting Public Keys

You can configure Email Security appliance to retrieve (harvest) public key from the incoming S/MIME signed messages and use it to send encrypted messages to the owner (business or consumer) of the harvested key.



#### Note

By default, public keys from expired or self-signed S/MIME certificates are not harvested.

### Before You Begin

Make sure that the public key of the sender's S/MIME certificate meets the requirements described in [S/MIME Certificate Requirements, page 19-20](#).

**Procedure**

- 
- Step 1** Click **Mail Policies > Mail Flow Policies**.
- Step 2** Create a new Mail Flow Policy or modify an existing one. See [Defining Which Hosts Are Allowed to Connect Using the Host Access Table \(HAT\)](#), page 7-1.
- Step 3** Scroll down to the **Security Features** section.
- Step 4** Under S/MIME Public Key Harvesting, do the following:
- Enable S/MIME public key harvesting.
  - (Optional) Choose whether to harvest public keys if the verification of the incoming signed messages fail.
  - (Optional) Choose whether to harvest updated public keys.



**Note** If an appliance receives more than one updated public key from the same domain or message within 48 hours, it sends out a warning alert.

---

- Step 5** Submit and commit your changes.
- 

**Note**

The size of the harvested public key repository on the appliance is 512 MB. If repository is full, Email Security appliance will automatically remove unused public keys.

---

**Note**

Use the `listenerconfig` command to enable key harvesting using CLI.

---

**Next Step**

Request the recipient to send a signed message to the Email Security appliance administrator. The Email Security appliance will harvest the public key from the signed message and displays it on the Mail Policies > Harvested Public Keys page.

## Managing S/MIME Sending Profiles

An S/MIME sending profile allows you define parameters such as:

- S/MIME mode to use, for example, sign, encrypt, and so on.
- S/MIME certificate for signing
- S/MIME signing mode to use, for example, opaque or detached.
- Action to take if the public key of the recipient's S/MIME certificate is not available on the appliance.

For example, one organization requires all the messages sent to them be signed and another one requires all the messages sent to them be signed and encrypted. In this scenario, you must create two sending profiles, one for signing alone and one for signing and encryption.

You can create, edit, delete, import, export, and search S/MIME sending profiles using the web interface or CLI.

## Create an S/MIME Sending Profile for Signing, Encrypting, or Signing and Encrypting Messages

### Procedure

**Step 1** Click **Mail Policies > Sending Profiles**.

**Step 2** Click **Add Profile**.

**Step 3** Configure the following fields:

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>S/MIME Profile Name</b> | Enter the name of the sending profile.                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>S/MIME Mode</b>         | <p>Choose the S/MIME mode. Possible values are:</p> <ul style="list-style-type: none"><li>• <b>Sign</b></li><li>• <b>Encrypt</b></li><li>• <b>Sign/Encrypt</b>. Sign and then encrypt</li><li>• <b>Triple</b>. Sign, encrypt, and then sign again</li></ul> <p><b>Note</b> If you are using one of the following S/MIME modes: <b>Sign</b>, <b>Sign/Encrypt</b>, or <b>Triple</b>, messages will be bounced to the sender if the signing fails.</p> |
| <b>Signing Certificate</b> | <p>Choose the signing certificate to use.</p> <p><b>Note</b> You need to set this field only if you choose one of the following S/MIME modes: <b>Sign</b>, <b>Sign/Encrypt</b>, or <b>Triple</b>.</p>                                                                                                                                                                                                                                               |

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>S/MIME Sign Mode</b> | <p>Choose the mode of S/MIME signing. Possible values are:</p> <ul style="list-style-type: none"> <li>• <b>Opaque.</b> An opaque-signed message contains the message and signature combined in a single part and can be read only by verifying the signature.</li> <li>• <b>Detached.</b> The signature information is separate from the text being signed. The MIME type for this is multipart/signed with the second part having a MIME subtype of application/(x-)pkcs7-signature.</li> </ul> <p><b>Note</b> You need to set this field only if you choose one of the following S/MIME modes: <b>Sign</b>, <b>Sign/Encrypt</b>, or <b>Triple</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>S/MIME Action</b>    | <p>Choose the action that Email Security appliance must take if the recipient's public key is not available. Possible values are:</p> <ul style="list-style-type: none"> <li>• <b>Bounce.</b> The message is bounced to the sender if any one of the recipient's public key is not available.</li> <li>• <b>Drop.</b> The message is dropped if any one of the recipient's public key is not available.</li> <li>• <b>Split.</b> The message is split. The message to the recipients whose public keys are not available are delivered without encryption and the message to the recipients whose public keys are available are encrypted and delivered.</li> </ul> <p><b>Example:</b> Assume that you are sending a message to bob@example1.com and dave@example2.com and the public key of dave@example2.com is not available. In this scenario, if you have selected <b>Split</b>, Email Security appliance will:</p> <ul style="list-style-type: none"> <li>– Deliver the message to bob@example1.com after encrypting it.</li> <li>– Deliver the message to dave@example2.com without encrypting it.</li> </ul> <p><b>Note</b> You need to set this field only if you choose one of the following S/MIME modes: <b>Encrypt</b>, <b>Sign/Encrypt</b>, or <b>Triple</b>.</p> |

**Step 4** Submit and commit your changes.



**Note**

Use the `smimeconfig` command to create sending profiles using CLI.

## Edit an S/MIME Sending Profile

- Step 1** Click **Mail Policies > Sending Profiles**.
- Step 2** Click on the sending profile that you want to modify.
- Step 3** Edit the fields as described in [Create an S/MIME Sending Profile for Signing, Encrypting, or Signing and Encrypting Messages](#), page 19-11.
- Step 4** Submit and commit your changes.

## Determining Which Messages to Sign, Encrypt, or Sign and Encrypt

After you create a sending profile, you need to create an outgoing content filter that determines which email messages should be signed, encrypted, or signed and encrypted. The content filter scans outgoing email and determines if the message matches the conditions specified. Once the content filter determines a message matches the condition, the Email Security appliance signs, encrypts, or signs or encrypts the message.

### Related Topics

- [Filtering Messages Based on Content, page 11-16](#)

## Signing, Encrypting, or Signing and Encrypting and Immediately Delivering Messages using a Content Filter

### Before You Begin

Understand the concept of building conditions for content filters. See [How Content Filters Work, page 11-1](#).

### Procedure

- 
- |                |                                                                                                                                                                                                                                                                                    |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b>  | Go to <b>Mail Policies &gt; Outgoing Content Filters</b> .                                                                                                                                                                                                                         |
| <b>Step 2</b>  | In the Filters section, click <b>Add Filter</b> .                                                                                                                                                                                                                                  |
| <b>Step 3</b>  | In the Conditions section, click <b>Add Condition</b> .                                                                                                                                                                                                                            |
| <b>Step 4</b>  | Add a condition to filter the messages that you want to sign, encrypt, or sign and encrypt. For example, to encrypt sensitive material, you might add a condition that identifies messages containing particular words or phrases, such as “Confidential,” in the subject or body. |
| <b>Step 5</b>  | Click <b>OK</b> .                                                                                                                                                                                                                                                                  |
| <b>Step 6</b>  | In the Actions section, click <b>Add Action</b> .                                                                                                                                                                                                                                  |
| <b>Step 7</b>  | Select <b>S/MIME Sign/Encrypt (Final Action)</b> from the <b>Add Action</b> list.                                                                                                                                                                                                  |
| <b>Step 8</b>  | Select the sending profile to associate with the content filter.                                                                                                                                                                                                                   |
| <b>Step 9</b>  | Click <b>OK</b> .                                                                                                                                                                                                                                                                  |
| <b>Step 10</b> | Submit and commit your changes.                                                                                                                                                                                                                                                    |
- 

### What To Do Next

After you add the content filter, you need to add the filter to an outgoing mail policy. You may want to enable the content filter on the default policy, or you may choose to apply the filter to a specific mail policy, depending on your organization’s needs. For information about working with mail policies, see [Overview of Mail Policies, page 10-1](#).

## Signing, Encrypting, or Signing and Encrypting a Message upon Delivery using a Content Filter

Create a content filter to sign, encrypt, or sign and encrypt a message on delivery, which means that the message continues to the next stage of processing, and when all processing is complete, the message is signed, encrypted, or signed and encrypted, and delivered.

### Before You Begin

- Understand the concept of building conditions for content filters. See [Overview of Content Filters, page 11-1](#).

### Procedure

- 
- |                |                                                                                                                                                                                                                                                                                    |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b>  | Go to <b>Mail Policies &gt; Outgoing Content Filters</b> .                                                                                                                                                                                                                         |
| <b>Step 2</b>  | In the Filters section, click <b>Add Filter</b> .                                                                                                                                                                                                                                  |
| <b>Step 3</b>  | In the Conditions section, click <b>Add Condition</b> .                                                                                                                                                                                                                            |
| <b>Step 4</b>  | Add a condition to filter the messages that you want to sign, encrypt, or sign and encrypt. For example, to encrypt sensitive material, you might add a condition that identifies messages containing particular words or phrases, such as “Confidential,” in the subject or body. |
| <b>Step 5</b>  | Click <b>OK</b> .                                                                                                                                                                                                                                                                  |
| <b>Step 6</b>  | In the Actions section, click <b>Add Action</b> .                                                                                                                                                                                                                                  |
| <b>Step 7</b>  | Select <b>S/MIME Sign/Encrypt on Delivery</b> from the <b>Add Action</b> list.                                                                                                                                                                                                     |
| <b>Step 8</b>  | Select the sending profile to associate with the content filter.                                                                                                                                                                                                                   |
| <b>Step 9</b>  | Click <b>OK</b> .                                                                                                                                                                                                                                                                  |
| <b>Step 10</b> | Submit and commit your changes.                                                                                                                                                                                                                                                    |
- 

### What To Do Next

After you add the content filter, you need to add the filter to an outgoing mail policy. You may want to enable the content filter on the default policy, or you may choose to apply the filter to a specific mail policy, depending on your organization’s needs. For information about working with mail policies, see [Overview of Mail Policies, page 10-1](#).

## Verifying, Decrypting, or Decrypting and Verifying Incoming Messages using S/MIME

- [S/MIME Verification and Decryption Workflow in Email Security Appliance, page 19-15](#)
- [How to Verify, Decrypt, or Decrypt and Verify Incoming Messages Using S/MIME, page 19-16](#)
- [Setting Up Certificates for Decrypting Messages, page 19-16](#)
- [Setting Up Public Keys for Verifying Signed Messages, page 19-17](#)
- [Enabling S/MIME Decryption and Verification, page 19-19](#)
- [Configuring an Action for S/MIME Decrypted or Verified Message, page 19-20](#)

**Note**

You can use Email Security appliance S/MIME security services to verify, decrypt, or decrypt and verify outgoing and incoming messages.

## S/MIME Verification and Decryption Workflow in Email Security Appliance

- [S/MIME Verification Workflow, page 19-15](#)
- [S/MIME Decryption Workflow, page 19-15](#)

### S/MIME Verification Workflow

The following process describes how Email Security appliance performs S/MIME verification.

1. Apply a hash algorithm to the signed message to create a message digest.
2. Decrypt the PKCS7 signature attached to the signed message using the public key of the sender's S/MIME certificate, and get the message digest.
3. Compare the generated message digest with the message digest retrieved from the signed message. If the message digests match, the message is verified.

### S/MIME Decryption Workflow

The following process describes how Email Security appliance performs S/MIME decryption.

1. Decrypt the session key using the private key of the appliance's S/MIME certificate
2. Decrypt the message body using the session key.



## How to Verify, Decrypt, or Decrypt and Verify Incoming Messages Using S/MIME

| Steps  | Do This                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | More Info                                                                                                                                                                                                                   |
|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Understand the S/MIME certificate requirements.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | See <a href="#">S/MIME Certificate Requirements, page 19-20</a> .                                                                                                                                                           |
| Step 2 | Depending on your requirements, do one of the following: <ul style="list-style-type: none"> <li>For S/MIME decryption, add your organization's S/MIME certificate (that contains the private key required to perform decryption) to the appliance.</li> <li>For S/MIME verification, add the public key of the sender's S/MIME certificate required to perform verification to the appliance.</li> <li>For S/MIME decryption and verification, add the following to the appliance:               <ul style="list-style-type: none"> <li>Your organization's S/MIME certificate (that contains the private key required to perform decryption) to the appliance.</li> <li>Public key of the sender's S/MIME certificate required to perform verification.</li> </ul> </li> </ul> | See <ul style="list-style-type: none"> <li><a href="#">Setting Up Certificates for Decrypting Messages, page 19-16</a></li> <li><a href="#">Setting Up Public Keys for Verifying Signed Messages, page 19-17</a></li> </ul> |
| Step 3 | Configure your mail flow policies to verify, decrypt, or decrypt and verify incoming messages using S/MIME.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | See <a href="#">Enabling S/MIME Decryption and Verification, page 19-19</a> .                                                                                                                                               |
| Step 4 | (Optional) Define the action that the Email Security appliance takes on decrypted or verified messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | See <a href="#">Configuring an Action for S/MIME Decrypted or Verified Message, page 19-20</a> .                                                                                                                            |



### Note

If you want to perform S/MIME verification, decryption, or decryption and verification using CLI, use the `listenerconfig > hostaccess` command. See the CLI inline help for more details.

## Setting Up Certificates for Decrypting Messages

You must add your organization's S/MIME certificate (that contains the private key required to perform decryption) to the appliance.

**Before You Begin**

- Share the public key of the appliance's S/MIME certificate with the sender (business or consumer) in one of the following ways:
  - Send the public key using an electronic channels, for example, email.
  - Request the sender to retrieve the public key using key harvesting.

The sender can use this public key to send encrypted messages to your appliance.

**Note**

In a B2C scenario, if your organization's S/MIME certificate is a domain certificate, some email clients (for example, Microsoft Outlook) may not be able to send encrypted messages using the public key of your organization's S/MIME certificate. This is because these email clients do not support encryption using public keys of domain certificates.

- Make sure that the S/MIME certificate that you plan to import meets the requirements described in [S/MIME Certificate Requirements, page 19-20](#).

**Procedure**

- 
- |               |                                                                          |
|---------------|--------------------------------------------------------------------------|
| <b>Step 1</b> | Click <b>Network &gt; Certificates</b> .                                 |
| <b>Step 2</b> | Click <b>Add Certificate</b> .                                           |
| <b>Step 3</b> | Choose <b>Import Certificate</b> .                                       |
| <b>Step 4</b> | Enter the path to the certificate file on your network or local machine. |
| <b>Step 5</b> | Enter the password for the file.                                         |
| <b>Step 6</b> | Click <b>Next</b> to view the certificate's information.                 |
| <b>Step 7</b> | Enter a name for the certificate.                                        |
| <b>Step 8</b> | Submit and commit your changes.                                          |
- 

**Note**

Use the `certconfig` command to add the S/MIME certificates using CLI.

## Setting Up Public Keys for Verifying Signed Messages

You must add the public key of the sender's S/MIME certificate to the appliance for verifying signed messages. Depending on your organizational policies and processes, you can use one of the following methods to add the public key to the appliance:

- Request the sender to send their public key using an electronic channels, for example, email. You can then add the public key using the web interface or CLI.

For instructions to add the public key, see [Adding a Public Key for S/MIME Verification, page 19-18](#).

- Retrieve the public key using key harvesting. See [Harvesting Public Keys for S/MIME Verification, page 19-18](#).

## Adding a Public Key for S/MIME Verification

### Before You Begin

- Make sure that the public key meets the requirements described in [S/MIME Certificate Requirements](#), page 19-20.
- Make sure that the public key is in PEM format.

### Procedure

- 
- |               |                                               |
|---------------|-----------------------------------------------|
| <b>Step 1</b> | Click <b>Mail Policies &gt; Public Keys</b> . |
| <b>Step 2</b> | Click <b>Add Public Key</b> .                 |
| <b>Step 3</b> | Enter the name of the public key.             |
| <b>Step 4</b> | Enter the public key.                         |
| <b>Step 5</b> | Submit and commit your changes.               |
- 



#### Note

Use the `smimeconfig` command to add public keys using CLI.

## Harvesting Public Keys for S/MIME Verification

You can configure Email Security appliance to retrieve (harvest) public key from the incoming S/MIME signed messages and use it to verify signed messages from the owner (business or consumer) of the harvested key.



#### Note

By default, public keys from expired or self-signed S/MIME certificates are not harvested.

### Procedure

1. Enable public key harvesting using the web interface or CLI. See [Enabling Public Key Harvesting](#), page 19-18.
2. Request the sender to send a signed message.
3. After the harvesting is complete, add the harvested public key to the appliance. See [Adding a Harvested Public Key for S/MIME Verification](#), page 19-19.

This step is to ensure that the message is verified at the gateway level.

## Enabling Public Key Harvesting

### Procedure

- 
- |               |                                                                                                                                                                        |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Click <b>Mail Policies &gt; Mail Flow Policies</b> .                                                                                                                   |
| <b>Step 2</b> | Create a new Mail Flow Policy or modify an existing one. See <a href="#">Defining Which Hosts Are Allowed to Connect Using the Host Access Table (HAT)</a> , page 7-1. |

**Step 3** Scroll down to the **Security Features** section.

**Step 4** Under S/MIME Public Key Harvesting, do the following:

- Enable S/MIME public key harvesting.
- (Optional) Choose whether to harvest public keys if the verification of the incoming signed messages fail.
- (Optional) Choose whether to harvest updated public keys.



**Note**

If an appliance receives more than one updated public key from the same domain or message within 48 hours, it sends out a warning alert.

**Step 5** Submit and commit your changes.



**Note**

The size of the harvested public key repository on the appliance is 512 MB. If the repository is full used, Email Security appliance will automatically remove unused public keys.



**Note**

Use the `listenerconfig` command to enable key harvesting using CLI.

## Adding a Harvested Public Key for S/MIME Verification

### Procedure

**Step 1** Click **Mail Policies > Harvested Public Keys**.

**Step 2** Click on the intended harvested public key and copy the public key.

**Step 3** Add the public key to the appliance. See [Adding a Public Key for S/MIME Verification, page 19-18](#).

**Step 4** Submit and commit your changes.

## Enabling S/MIME Decryption and Verification

### Procedure

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

**Step 2** Create a new Mail Flow Policy or modify an existing one. See [Defining Which Hosts Are Allowed to Connect Using the Host Access Table \(HAT\), page 7-1](#).

**Step 3** Scroll down to the **Security Features** section.

**Step 4** Under S/MIME Decryption/Verification, do the following:

- Enable S/MIME decryption and verification.

- Choose whether to retain or remove the digital signature from the messages after S/MIME verification. If you do not want your end users to know about S/MIME gateway verification, select **Remove**.

For triple wrapped messages, only the inner signature is retained or removed.

**Step 5** Submit and commit your changes.

**Tip**

If S/MIME Decryption and Verification is enabled in the Mail Flow Policies, all the S/MIME messages are delivered irrespective of the status of the decryption and verification. If you want to configure an action for handling S/MIME Decrypted or Verified Messages, you can use the message filter rules—`smime-gateway-verified` and `smime-gateway`. For more information, see [Configuring an Action for S/MIME Decrypted or Verified Message](#), page 19-20.

## Configuring an Action for S/MIME Decrypted or Verified Message

After Email Security appliance performs S/MIME decryption, verification, or both, you may want to take different actions depending on the results. You can use the message filter rules—`smime-gateway-verified` and `smime-gateway` to perform actions on the messages based on the result of decryption, verification, or both. For more information, see [Chapter 9, “Using Message Filters to Enforce Email Policies.”](#)

**Note**

You can also use the content filter conditions—**S/MIME Gateway Message** and **S/MIME Gateway Verified** to perform actions on the messages based on the result of decryption, verification, or both. For more information, see [Chapter 11, “Content Filters.”](#)

### Example: Quarantine S/MIME Messages that failed Verification, Decryption, or Both

The following message filter checks if the message is an S/MIME message and quarantines it if the verification or decryption using S/MIME fails.

```
quarantine_smime_messages:if (smime-gateway-message and not smime-gateway-verified) {
 quarantine("Policy"); }
```

## S/MIME Certificate Requirements

- [Certificate Requirements for Signing](#), page 19-20
- [Certificate Requirements for Encryption](#), page 19-21

## Certificate Requirements for Signing

The S/MIME certificate for signing must contain the following information:

|                            |                                           |
|----------------------------|-------------------------------------------|
| <b>Common Name</b>         | The fully qualified domain name.          |
| <b>Organization</b>        | The exact legal name of the organization. |
| <b>Organizational Unit</b> | Section of the organization.              |

|                                          |                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>City (Locality)</b>                   | The city where the organization is legally located.                                                                                                                                                                                                                                                                                                                                                   |
| <b>State (Province)</b>                  | The state, county, or region where the organization is legally located.                                                                                                                                                                                                                                                                                                                               |
| <b>Country</b>                           | The two letter ISO abbreviation of the country where the organization is legally located.                                                                                                                                                                                                                                                                                                             |
| <b>Duration before expiration</b>        | The number of days before the certificate expires.                                                                                                                                                                                                                                                                                                                                                    |
| <b>Subject Alternative Name(Domains)</b> | Name of the domain from which you plan to send signed messages. Examples include <code>domain.com</code> and <code>*.domain.net</code> . For multiple entries, use a comma-separated list.                                                                                                                                                                                                            |
| <b>Subject Alternative Name(Email)</b>   | Email address of the user who is planning to send signed messages, for example, <code>user@somedomain.com</code> . For multiple entries, use a comma-separated list.                                                                                                                                                                                                                                  |
| <b>Private Key Size</b>                  | Size of the private key to generate for the CSR.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Key Usage</b>                         | Key usage is a restriction method that determines what a certificate can be used for. If the key usage extension is specified, the following bits: <code>digitalSignature</code> and <code>nonRepudiation</code> must be set.<br><br>If the key usage extension is not specified, receiving clients must presume that the <code>digitalSignature</code> and <code>nonRepudiation</code> bits are set. |

For detailed information about S/MIME certificates, see RFC 5750: Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 - Certificate Handling.

## Certificate Requirements for Encryption

The S/MIME certificate for encryption must contain the following information:

|                                          |                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Common Name</b>                       | The fully qualified domain name.                                                                                                                                                                                                                                                                                    |
| <b>Organization</b>                      | The exact legal name of the organization.                                                                                                                                                                                                                                                                           |
| <b>Organizational Unit</b>               | Section of the organization.                                                                                                                                                                                                                                                                                        |
| <b>City (Locality)</b>                   | The city where the organization is legally located.                                                                                                                                                                                                                                                                 |
| <b>State (Province)</b>                  | The state, county, or region where the organization is legally located.                                                                                                                                                                                                                                             |
| <b>Country</b>                           | The two letter ISO abbreviation of the country where the organization is legally located.                                                                                                                                                                                                                           |
| <b>Duration before expiration</b>        | The number of days before the certificate expires.                                                                                                                                                                                                                                                                  |
| <b>Subject Alternative Name(Domains)</b> | Name of the domain to which you plan to send encrypted messages. Examples include <code>domain.com</code> and <code>*.domain.net</code> . For multiple entries, use a comma-separated list.<br><br>If you plan to send encrypted messages to all the users in a domain, the public key should include a SAN Domain. |
| <b>Subject Alternative Name(Email)</b>   | Email address of the user to whom you plan to send encrypted messages, for example, <code>user@somedomain.com</code> . For multiple entries, use a comma-separated list.                                                                                                                                            |

|                         |                                                                                                                                                                                                   |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Private Key Size</b> | Size of the private key to generate for the CSR.                                                                                                                                                  |
| <b>Key Usage</b>        | Key usage is a restriction method that determines what a certificate can be used for. The key usage extension must be specified and the following bit must be set: <code>keyEncipherment</code> . |

For detailed information about S/MIME certificates, see RFC 5750: Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 - Certificate Handling.

#### Before You Begin

- Make sure that the public key meets the requirements described in [S/MIME Certificate Requirements, page 19-20](#).
- Make sure that the public key is in PEM format.

#### Procedure

- 
- Step 1** Click **Mail Policies > Public Keys**.
- Step 2** Click **Add Public Key**.
- Step 3** Enter the name of the public key.
- Step 4** Enter the public key.
- Step 5** Submit and commit your changes.
- 



#### Note

Use the `smimeconfig` command to add public keys using CLI.

---

#### Before You Begin

Copy the export file to the `/configuration` directory of the appliance. For instructions to create an export file, see [Exporting Public Keys, page 19-23](#).

#### Procedure

- 
- Step 1** Click **Mail Policies > Public Keys**.
- Step 2** Click **Import Public Keys**.
- Step 3** Select the export file and click **Submit**.



#### Note

The import process may take longer if you are importing a file with large number of public keys. Make sure that you adjust the web interface or CLI inactivity timeout accordingly.

---

- Step 4** Commit your changes.
-

## Exporting Public Keys

All public keys on the appliance are exported together in a single text file and stored in the `/configuration` directory.

### Procedure

- 
- |               |                                                     |
|---------------|-----------------------------------------------------|
| <b>Step 1</b> | Choose <b>Mail Policies &gt; Public Keys</b> .      |
| <b>Step 2</b> | Click <b>Export Public Keys</b> .                   |
| <b>Step 3</b> | Enter a name for the file and click <b>Submit</b> . |
-







# Email Authentication

---

- [Email Authentication Overview, page 20-1](#)
- [DomainKeys and DKIM Authentication, page 20-1](#)
- [Configuring DomainKeys and DKIM Signing, page 20-3](#)
- [How to Verify Incoming Messages Using DKIM, page 20-16](#)
- [Overview of SPF and SIDF Verification, page 20-22](#)
- [How to Verify Incoming Messages Using SPF/SIDF, page 20-23](#)
- [Enabling SPF and SIDF, page 20-24](#)
- [Determining the Action to Take for SPF/SIDF Verified Mail, page 20-31](#)
- [Testing the SPF/SIDF Results, page 20-34](#)
- [DMARC Verification, page 20-35](#)

## Email Authentication Overview

AsyncOS for Email supports email verification and signing to prevent email forgery. To verify incoming mail, AsyncOS supports Sender Policy Framework (SPF), Sender ID Framework (SIDF), DomainKeys Identified Mail (DKIM), and Domain-based Message Authentication, Reporting and Conformance (DMARC). To authenticate outbound mail, AsyncOS supports DomainKeys and DKIM signing.

### Related Topics

- [DomainKeys and DKIM Authentication, page 20-1.](#)
- [Overview of SPF and SIDF Verification, page 20-22.](#)
- [DMARC Verification, page 20-35](#)

## DomainKeys and DKIM Authentication

With DomainKeys or DKIM email authentication, the sender signs the email using public key cryptography. The verified domain can then be used to detect forgeries by comparing it with the domain in the From: (or Sender:) header of the email.

DomainKeys and DKIM consist of two main parts: signing and verification. AsyncOS 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.

**Related Topics**

- [DomainKeys and DKIM Authentication Workflow, page 20-2](#)
- [DomainKeys and DKIM Signing in AsyncOS, page 20-2](#)

## DomainKeys and DKIM Authentication Workflow

**Figure 20-1 Authentication Work Flow**



1. Administrator (domain owner) publishes a public key into the DNS name space.
2. Administrator loads a private key in the outbound Mail Transfer Agent (MTA).
3. Email submitted by an authorized user of that domain is digitally signed with the respective private key. The signature is inserted in the email as a DomainKey or DKIM signature header and the email is transmitted.
4. Receiving MTA extracts the DomainKeys or DKIM signature from the header and the claimed sending domain (via the Sender: or From: header) from the email. The public key is retrieved from the claimed signing domain which is extracted from DomainKeys or DKIM signature header fields.
5. The public key is used to determine whether the DomainKeys or DKIM signature was generated with the appropriate private key.

To test your outgoing DomainKeys signatures, you can use a Yahoo! or Gmail address, as these services are free and provide validation on incoming messages that are DomainKeys signed.

## DomainKeys and DKIM Signing in AsyncOS

DomainKeys and DKIM signing in AsyncOS is implemented via domain profiles and enabled via a mail flow policy (typically, the outgoing “relay” policy). For more information, see the “Configuring the Gateway to Receive Mail” chapter. Signing the message is the last action performed by the appliance before the message is sent.

Domain profiles associate a domain with domain key information (signing key and related information). As email is sent via a mail flow policy on the appliance, sender email addresses that match any domain profile are DomainKeys signed with the signing key specified in the domain profile. If you enable both DKIM and DomainKeys signing, the DKIM signature is used. You implement DomainKeys and DKIM profiles via the `domainkeysconfig` CLI command or via the Mail Policies > Domain Profiles and the Mail Policies > Signing Keys pages in the GUI.

DomainKeys and DKIM signing works like this: a domain owner generates two keys — a public key stored in the public DNS (a DNS TXT record associated with that domain) and a private key that is stored on the appliance is used to sign mail that is sent (mail that originates) from that domain.

As messages are received on a listener used to send messages (outbound), the appliance checks to see if any domain profiles exist. If there are domain profiles created on the appliance (and implemented for the mail flow policy), the message is scanned for a valid Sender: or From: address. If both are present, the Sender: is used for DomainKeys. The From: address is always used for DKIM signing. Otherwise, the first From: address is used. If a valid address is not found, the message is not signed and the event is logged in the mail\_logs.

**Note**

If you create both a DomainKey and DKIM profile (and enable signing on a mail flow policy), AsyncOS signs outgoing messages with both a DomainKeys and DKIM signature.

If a valid sending address is found, the sending address is matched against the existing domain profiles. If a match is found, the message is signed. If not, the message is sent without signing. If the message has an existing DomainKeys (a "DomainKey-Signature:" header) the message is only signed if a new sender address has been added after the original signing. If a message has an existing DKIM signature, a new DKIM signature is added to the message.

AsyncOS provides a mechanism for signing email based on domain as well as a way to manage (create new or input existing) signing keys.

The configuration descriptions in this document represent the most common uses for signing and verification. You can also enable DomainKeys and DKIM signing on a mail flow policy for inbound email, or enable DKIM verification on a mail flow policy for outbound email.

**Note**

When you configure domain profiles and signing keys in a clustered environment, note that the Domain Key Profile settings and Signing Key settings are linked. Therefore, if you copy, move or delete a signing key, the same action is taken on the related profile.

## Configuring DomainKeys and DKIM Signing

### Related Topics

- [Signing Keys, page 20-3](#)
- [Public Keys, page 20-4](#)
- [Domain Profiles, page 20-5](#)
- [Enabling Signing for Outgoing Mail, page 20-6](#)
- [Enabling Signing for Bounce and Delay Messages, page 20-6](#)
- [Configuring DomainKeys/DKIM Signing \(GUI\), page 20-7](#)
- [Domain Keys and Logging, page 20-16](#)

## Signing Keys

A signing key is the private key stored on the appliance. When creating a signing key, you specify a key size. Larger key sizes are more secure; however, larger keys also can impact performance. The appliance supports keys from 512 bits up to 2048 bits. The 768 - 1024 bit key sizes are considered secure and used by most senders today. Keys based on larger key sizes can impact performance and are not supported above 2048 bits. For more information about creating signing keys, see [Creating or Editing a Signing Key, page 20-10](#).

If you are entering an existing key, simply paste it into the form. Another way to use existing signing keys is to import the key as a text file. For more information about adding existing signing keys, see [Importing or Entering Existing Signing Keys, page 20-11](#).

Once a key is entered, it is available for use in domain profiles, and will appear in the Signing Key drop-down list in the domain profile.

#### Related Topics

- [Exporting and Importing Signing Keys, page 20-4](#)

## Exporting and Importing Signing Keys

You can export your signing keys to a text file on the appliance. When you export keys, all of the keys currently existing on the appliance are put into a text file. For more information about exporting keys, see [Exporting Signing Keys, page 20-11](#).

You can import keys that have been exported as well.



#### Note

Importing keys causes all of the current keys on the appliance to be replaced.

For more information, see [Importing or Entering Existing Signing Keys, page 20-11](#).

## Public Keys

Once you have associated a signing key with a domain profile, you can create DNS text record which contains your public key. You do this via the Generate link in the DNS Text Record column in the domain profile listing (or via `domainkeysconfig -> profiles -> dnstxt` in the CLI):

**Figure 20-2** Generate DNS Text Record Link on Domain Profiles Page

| Domain Profiles                    |             |          |              |                                    |                          |                                    |                          |
|------------------------------------|-------------|----------|--------------|------------------------------------|--------------------------|------------------------------------|--------------------------|
| <a href="#">Add Profile...</a>     |             |          |              | <a href="#">Clear All Profiles</a> |                          | <a href="#">Import Profiles...</a> |                          |
| Profile Name                       | Domain      | Selector | Users        | Signing Key                        | DNS Text Record          | Test Profile                       | All<br>Delete            |
| ExampleProfile                     | example.com | test     | .example.com | myTestKey                          | <a href="#">Generate</a> | Test                               | <input type="checkbox"/> |
| <a href="#">Export Profiles...</a> |             |          |              | <a href="#">Delete</a>             |                          |                                    |                          |

For more information about generating a DNS Text Record, see [Generating a DNS Text Record, page 20-13](#).

You can also view the public key via the View link on the Signing Keys page:

**Figure 20-3** View Public Key Link on Signing Keys Page  
**Signing Keys**

| Signing Keys                   |                 |                                |                 |                                |
|--------------------------------|-----------------|--------------------------------|-----------------|--------------------------------|
| <a href="#">Add Key...</a>     |                 | <a href="#">Clear All Keys</a> |                 | <a href="#">Import Keys...</a> |
| Name                           | Key Size (Bits) | Public Key                     | Domain Profiles | All<br>Delete                  |
| TestKey                        | 768             | <a href="#">View</a>           | ExampleProfile  | <input type="checkbox"/>       |
| <a href="#">Export Keys...</a> |                 |                                |                 | <a href="#">Delete</a>         |

## Domain Profiles

A domain profile associates a sender domain with a signing key, along with some other information needed for signing.

- A name for the domain profile.
- A domain name (the domain to be included in the “d=” header).
- A selector (a selector is used to form the query for the public key. In the DNS query type, this value is prepended to the “\_domainkey.” namespace of the sending domain).
- A canonicalization method (the method by which the headers and content are prepared for presentation to the signing algorithm). AsyncOS supports both “simple” and “nofws” for DomainKeys and “relaxed” and “simple” for DKIM.
- A signing key (see [Signing Keys, page 20-3](#) for more information).
- A list of headers and the body length to sign (DKIM only).
- A list of tags you want to include in the signature’s header (DKIM only). These tags store the following information:
  - The identity of the user or agent (e.g., a mailing list manager) on whose behalf the message is signed.
  - A comma-separated list of query methods used to retrieve the public key.
  - The timestamp of when the signature was created.
  - The expiration time of the signature, in seconds.
  - A vertical bar-separated (i.e., |) list of header fields present when the message was signed.
- The tags you want to include in the signature (DKIM only).
- A list of Profile Users (addresses allowed to use the domain profile for signing).



### Note

The domain in the addresses specified in the profile users must match the domain specified in the Domain field.

You can search through all of your existing domain profiles for a specific term. See [Searching Domain Profiles, page 20-15](#) for more information.

You can also choose whether to sign system-generated messages with DKIM signatures. See [Signing System-Generated Messages, page 20-15](#) for more information.

### Related Topics

- [Exporting and Importing Domain Profiles, page 20-6](#)

## Exporting and Importing Domain Profiles

You can export your existing domain profiles to a text file on the appliance. When you export the domain profiles, all of the profiles existing on the appliance are put into a single text file. See [Exporting Domain Profiles, page 20-14](#).

You can import domain profiles that you have previously exported. Importing domain profiles causes all of the current domain profiles on the machine to be replaced. See [Importing Domain Profiles, page 20-14](#).

## Enabling Signing for Outgoing Mail

DomainKeys and DKIM signing is enabled on mail flow policies for outbound mail. For more information, see the “Configuring the Gateway to Receive Mail” chapter.

### Procedure

- 
- Step 1** On the Mail Flow Policies page (from the Mail Policies menu), click on the RELAYED mail flow policy (outgoing).
  - Step 2** From the Security Features section, enable DomainKeys/DKIM Signing by selecting On.
  - Step 3** Submit and commit your changes.
- 

## Enabling Signing for Bounce and Delay Messages

In addition to signing outbound messages, you may want to sign bounce and delay messages. You may want to do this to alert recipients that the bounce and delay messages they receive from your company are legitimate. To enable DomainKeys and DKIM signing for bounce and delay messages, you enable DomainKeys/DKIM signing for the bounce profile associated with the public listener.

### Procedure

- 
- Step 1** On the bounce profile associated with the public listener where you will send signed outbound messages, go to Hard Bounce and Delay Warning Messages.
  - Step 2** Enable “Use Domain Key Signing for Bounce and Delay Messages”:



#### Note

You must have completed all steps listed in [Configuring DomainKeys/DKIM Signing \(GUI\), page 20-7](#) to sign bounced and delay messages.



#### Note

The From: address in the domain profile must match the address used for the bounce return address. To ensure these addresses match, you can configure a return address for the bounce profile (System Administration > Return Addresses), and then use the same name in the Profile Users list in the domain

profile. For example, you would configure a return address of MAILER-DAEMON@example.com for the bounce return address, and add MAILER-DAEMON@example.com as a profile user in the domain profile.

---

## Configuring DomainKeys/DKIM Signing (GUI)

### Procedure

---

- Step 1** Create a new or import an existing private key. For information on creating or importing signing keys, see [Signing Keys, page 20-3](#).
- Step 2** Create a domain profile and associate the key with the domain profile. For information on creating a domain profile, see [Domain Profiles, page 20-5](#).
- Step 3** Create the DNS text record. For information about creating the DNS text record, see [Generating a DNS Text Record, page 20-13](#).
- Step 4** If you have not already done so, enable DomainKeys/DKIM signing on a mail flow policy for outbound mail (see [Enabling Signing for Outgoing Mail, page 20-6](#)).
- Step 5** Optionally, enable DomainKeys/DKIM signing for bounced and delay messages. For information about enabling signing for bounce and delay messages, see [Enabling Signing for Bounce and Delay Messages, page 20-6](#).
- Step 6** Send email. Mail sent from a domain that matches a domain profile will be DomainKeys/DKIM signed. In addition, bounce or delay messages will be signed if you configured signing for bounce and delay messages.



### Note

If you create both a DomainKey and DKIM profile (and enable signing on a mail flow policy), AsyncOS signs outgoing messages with both a DomainKeys and DKIM signature.

---

### Related Topics

- [Creating Domain Profiles for DomainKeys Signing, page 20-8](#)
- [Creating a New Domain Profile for DKIM Signing, page 20-8](#)
- [Creating or Editing a Signing Key, page 20-10](#)
- [Exporting Signing Keys, page 20-11](#)
- [Importing or Entering Existing Signing Keys, page 20-11](#)
- [Deleting Signing Keys, page 20-12](#)
- [Generating a DNS Text Record, page 20-13](#)
- [Testing Domain Profiles, page 20-14](#)
- [Exporting Domain Profiles, page 20-14](#)
- [Importing Domain Profiles, page 20-14](#)
- [Deleting Domain Profiles, page 20-14](#)



- [Searching Domain Profiles, page 20-15](#)
- [Signing System-Generated Messages, page 20-15](#)

## Creating Domain Profiles for DomainKeys Signing

### Procedure

- 
- Step 1** Choose **Mail Policies > Signing Profiles**.
- Step 2** In the **Domain Signing Profiles** section, click **Add Profile**.
- Step 3** Enter a name for the profile.
- Step 4** For the **Domain Key Type**, choose **Domain Keys**.  
Additional options appear on the page.
- Step 5** Enter the domain name.
- Step 6** Enter a selector. Selectors are arbitrary names prepended to the "\_domainkey" namespace, used to help support multiple concurrent public keys per sending domain. A selector value and length must be legal in the DNS namespace and in email headers with the additional provision that they cannot contain a semicolon.
- Step 7** Select the canonicalization (no forwarding whitespaces or simple).
- Step 8** If you have already created a signing key, select a signing key. Otherwise, skip to the next step. You must create (or import) at least one signing key in order to have signing keys to choose from in the list. See [Creating or Editing a Signing Key, page 20-10](#).
- Step 9** Enter users (email addresses, hosts, etc.) that will use the domain profile for signing.
- Step 10** Submit and commit your changes.
- Step 11** At this point (if you have not already) you should enable DomainKeys/DKIM signing on an outgoing mail flow policy (see [Enabling Signing for Outgoing Mail, page 20-6](#)).



---

**Note** If you create both a DomainKeys and DKIM profile, AsyncOS performs both DomainKeys and DKIM signing on outgoing mail.

---

## Creating a New Domain Profile for DKIM Signing

### Procedure

- 
- Step 1** Choose **Mail Policies > Signing Profiles**.
- Step 2** In the **Domain Signing Profiles** section, click **Add Profile**.
- Step 3** Enter a name for the profile.
- Step 4** For the **Domain Key Type**, choose **DKIM**.  
Additional options appear on the page.
- Step 5** Enter the domain name.

- Step 6** Enter a selector. Selectors are arbitrary names prepended to the "\_domainkey." namespace, used to help support multiple concurrent public keys per sending domain. A selector value and length must be legal in the DNS namespace and in email headers with the additional provision that they cannot contain a semicolon.
- Step 7** Select the canonicalization for the header. Choose from the following options:
- **Relaxed.** The “relaxed” header canonicalization algorithm performs the following: header names are changed to lowercase, headers are unfolded, linear white spaces are reduced to a single space, leading and trailing spaces are stripped.
  - **Simple.** No changes to headers are made.
- Step 8** Select the canonicalization for the body. Choose from the following options:
- **Relaxed.** The “relaxed” header canonicalization algorithm performs the following: empty lines are stripped at the end of the body, white spaces are reduced to a single space within lines, and trailing white spaces are stripped in lines.
  - **Simple.** Empty lines at the end of the body are stripped.
- Step 9** If you have already created a signing key, select a signing key. Otherwise, skip to the next step. You must create (or import) at least one signing key in order to have signing keys to choose from in the list. See [Creating or Editing a Signing Key, page 20-10](#).
- Step 10** Select the list of headers to sign. You can select from the following headers:
- **All.** AsyncOS signs all the headers present at the time of signature. You may want to sign all headers if you do not expect headers to be added or removed in transit.
  - **Standard.** You may want to select the standard headers if you expect that headers may be added or removed in transit. AsyncOS signs only the following standard headers (if the header is not present in the message, the DKIM signature indicates a null value for the header):
    - From
    - Sender, Reply To-
    - Subject
    - Date, Message-ID
    - To, Cc
    - MIME-Version
    - Content-Type, Content-Transfer-Encoding, Content-ID, Content-Description
    - Resent-Date, Resent-From, Resent-Sender, Resent-To, Resent-cc, Resent-Message-ID
    - In-Reply-To, References
    - List-Id, List-Help, List-Unsubscribe, List-Subscribe, List-Post, List-Owner, List-Archive

**Note**

When you select “Standard”, you can add additional headers to sign.

- Step 11** Specify how to sign the message body. You can choose to sign the message body, and/or how many bytes to sign. Select one of the following options:
- **Whole Body Implied.** Do not use the “l=” tag to determine body length. The entire message is signed and no changes are allowed.
  - **Whole Body Auto-determined.** The entire message body is signed, and appending some additional data to the end of body is allowed during transit.

- **Sign first \_ bytes.** Sign the message body up to the specified number of bytes.

**Step 12** Select the tags you want to include in the message signature's header field. The information stored in these tags are used for message signature verification. Select one or more of the following options:

- **“i” Tag.** The identity of the user or agent (e.g., a mailing list manager) on behalf of which this message is signed. Enter the domain name prepended with the @ symbol, such as the domain `@example.com`.
- **“q” Tag.** A colon-separated list of query methods used to retrieve the public key. Currently, the only valid value is `dns/txt`.
- **“t” Tag.** A timestamp for when the signature was created.
- **“x” Tag.** The absolute date and time when the signature expires. Specify an expiration time (in seconds) for the signature. The default is `31536000` seconds.
- **“z” Tag.** A vertical bar-separated (i.e., |) list of header fields present when the message was signed. This includes the names of the header fields and their values. For example:

```
z=From:admin@example.com|To:joe@example.com|
Subject:test%20message|Date:Date:August%2026,%202011%205:30:02%20PM%20-0700
```

**Step 13** Enter users (email addresses, hosts, etc.) that will use the domain profile for signing.



**Note**

When you create domain profiles, be aware that a hierarchy is used in determining the profile to associate with a particular user. For example, you create a profile for `example.com` and another profile for `joe@example.com`. When mail is sent from `joe@example.com`, the profile for `joe@example.com` is used. However, when mail is sent from `adam@example.com`, the profile for `example.com` is used.

**Step 14** Submit and commit your changes.

**Step 15** At this point (if you have not already) you should enable DomainKeys/DKIM signing on an outgoing mail flow policy (see [Enabling Signing for Outgoing Mail, page 20-6](#)).



**Note**

If you create both a DomainKeys and DKIM profile, AsyncOS performs both DomainKeys and DKIM signing on outgoing mail.

## Creating or Editing a Signing Key

- [Create a New Signing Key, page 20-10](#)
- [Edit an Existing Signing Key, page 20-11](#)

### Create a New Signing Key

Signing keys are required for domain profiles for DomainKeys and DKIM signing.

#### Procedure

**Step 1** Choose **Mail Policies > Signing Keys**.

**Step 2** Click **Add Key**.

**Step 3** Enter a name for the key.

**Step 4** Click **Generate** and select a key size.

**Step 5** Submit and commit your changes.



**Note** If you have not done so already, you may need to edit your domain profile to assign the key.

## Edit an Existing Signing Key

### Procedure

**Step 1** Choose **Mail Policies > Signing Keys**.

**Step 2** Click the intended signing key.

**Step 3** Edit the intended fields as described in .



**Note** For enhanced security, if encryption of sensitive data in the appliance is enabled in FIPS mode, you will not be able view the private key. If you intend to edit the private key, you can paste your private key or generate a new private key.

**Step 4** Submit and commit your changes.

## Exporting Signing Keys

All keys on the appliance are exported together in a single text file.

### Procedure

**Step 1** Choose **Mail Policies > Signing Keys**.

**Step 2** Click **Export Keys**.



**Note** For enhanced security, if encryption of sensitive data in the appliance is enabled in FIPS mode, signing keys are encrypted while exporting.

**Step 3** Enter a name for the file and click **Submit**.

## Importing or Entering Existing Signing Keys

### Related Topics

- [Pasting a Key, page 20-12](#)
- [Importing Keys from an Existing Export File, page 20-12](#)

## Pasting a Key

### Procedure

---

- Step 1** Choose **Mail Policies > Signing Keys**.
- Step 2** Click **Add Key**.
- Step 3** Paste the key into the Paste Key field (must be PEM-formatted and must be RSA keys only).
- Step 4** Submit and commit your changes.
- 

## Importing Keys from an Existing Export File



### Note

To obtain a key file, see [Exporting Signing Keys, page 20-11](#).

---

### Procedure

---

- Step 1** Choose **Mail Policies > Signing Keys**.
- Step 2** Click **Import Keys**.
- Step 3** Select the file that contains the exported signing keys.
- Step 4** Click **Submit**. You are warned that importing will replace all existing signing keys. All of the keys in the text file are imported.
- Step 5** Click **Import**.
- 

## Deleting Signing Keys

### Related Topics

- [Removing Selected Signing Keys, page 20-12](#)
- [Removing All Signing Keys, page 20-13](#)

## Removing Selected Signing Keys

### Procedure

---

- Step 1** Choose **Mail Policies > Signing Keys**.
- Step 2** Mark the checkbox to the right of each signing key to remove.
- Step 3** Click **Delete**.
- Step 4** Confirm the deletion.
-

## Removing All Signing Keys

### Procedure

- 
- Step 1** Choose **Mail Policies > Signing Keys**.
- Step 2** Click **Clear All Keys** on the Signing Keys page.
- Step 3** Confirm the deletion.
- 

## Generating a DNS Text Record

### Procedure

- 
- Step 1** Choose **Mail Policies > Signing Profiles**.
- Step 2** In the Domain Signing Profiles section, in the DNS Text Record column, click the **Generate** link for the corresponding domain profile.
- Step 3** Mark the checkbox for the attributes you wish to include in the DNS text record.
- Step 4** Click **Generate Again** to re-generate the key with any changes you have made.
- Step 5** The DNS text record is displayed in the text field at the bottom of the window (where you can now copy it). In some cases, multi-string DNS text records are generated. See [Multi-string DNS Text Records, page 20-13](#).
- Step 6** Click **Done**.
- 

### Related Topics

- [Multi-string DNS Text Records, page 20-13](#)

## Multi-string DNS Text Records

Multi-string DNS text records may be generated if the key size of the signing key used to generate the DNS text records are larger than 1024 bits. This is because not more than 255 characters are allowed in a single string of a DNS text record. The DKIM authentication may fail as some of the DNS servers do not accept or serve multi-string DNS text records.

To avoid this scenario, it is recommended that you use double quotes to break up the multi-string DNS text record into smaller strings not exceeding 255 bytes. The following is an example.

```
s._domainkey.domain.com. IN TXT "v=DKIM1;"
"p=MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQE"
"A4Vbhjq2n/3DbEk6EHdeVXlIXFT7OE181amoZLbvWMX+bej"
"CdxcsFV3uS7G8oOJSWBP0z++nTQmy9ZDWfaiopU6k7tzoI"
"+oRDlKkhCQrM4oP2B2F5sTDkYwPY3Pen2jgC2OgbPnbo3o"
"m3c1wMWgSoZxoZUE4ly5kPuK9fTtpeJHNIzAqkFICiev4yrkL"
"R+SmFsJn9MYH5+1chyZ74BVm+16Xq2mptWXEwpiwOxWI"
"YHXsZo2zRjedrQ45vmgb8xUx5ioYY9/yBLHudGc+GUKTj1i4"
"mQg48yCD/HVNfSRXaPinliEkypH9cSnvgvWuIYUQz0dHU;"
```

DKIM implementations reassemble DNS text records broken down this way into the full original single string before processing them.

## Testing Domain Profiles

Once you have created a signing key, associated it with a domain profile, and generated and inserted the DNS text into your authorized DNS, you can test your domain profile.

### Procedure

- 
- |               |                                                                                                                                                            |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Choose <b>Mail Policies &gt; Signing Profiles</b> .                                                                                                        |
| <b>Step 2</b> | In the <b>Domain Signing Profiles</b> section, in the Test Profile column, click the <b>Test</b> link for the domain profile.                              |
| <b>Step 3</b> | A message is displayed at the top of the page, indicating success or failure. If the test fails, a warning message is displayed, including the error text. |
- 

## Exporting Domain Profiles

All domain profiles on the appliance are exported together in a single text file.

### Procedure

- 
- |               |                                                     |
|---------------|-----------------------------------------------------|
| <b>Step 1</b> | Choose <b>Mail Policies &gt; Signing Profiles</b> . |
| <b>Step 2</b> | Click <b>Export Domain Profiles</b> .               |
| <b>Step 3</b> | Enter a name for the file and click <b>Submit</b> . |
- 

## Importing Domain Profiles

### Procedure

- 
- |               |                                                                                                                                                          |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Choose <b>Mail Policies &gt; Signing Profiles</b> .                                                                                                      |
| <b>Step 2</b> | Click <b>Import Domain Profiles</b> .                                                                                                                    |
| <b>Step 3</b> | Select the file that contains the exported domain profiles.                                                                                              |
| <b>Step 4</b> | Click <b>Submit</b> . You are warned that importing will replace all existing domain profiles. All of the domain profiles in the text file are imported. |
| <b>Step 5</b> | Click <b>Import</b> .                                                                                                                                    |
- 

## Deleting Domain Profiles

### Related Topics

- [Removing Selected Domain Profiles, page 20-15](#)
- [Removing All Domain Profiles, page 20-15](#)

## Removing Selected Domain Profiles

### Procedure

- 
- |               |                                                                  |
|---------------|------------------------------------------------------------------|
| <b>Step 1</b> | Choose <b>Mail Policies &gt; Signing Profiles</b> .              |
| <b>Step 2</b> | Mark the checkbox to the right of each domain profile to remove. |
| <b>Step 3</b> | Click <b>Delete</b> .                                            |
| <b>Step 4</b> | Confirm the deletion.                                            |
- 

## Removing All Domain Profiles

### Procedure

- 
- |               |                                                     |
|---------------|-----------------------------------------------------|
| <b>Step 1</b> | Choose <b>Mail Policies &gt; Signing Profiles</b> . |
| <b>Step 2</b> | Click <b>Clear All Profiles</b> .                   |
| <b>Step 3</b> | Confirm the deletion.                               |
- 

## Searching Domain Profiles

### Procedure

- 
- |               |                                                                                                               |
|---------------|---------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Choose <b>Mail Policies &gt; Signing Profiles</b> .                                                           |
| <b>Step 2</b> | In the Find Domain Profiles section, specify the search term.                                                 |
| <b>Step 3</b> | Click <b>Find Profiles</b> .                                                                                  |
| <b>Step 4</b> | The search scans the following fields for each domain profile: email, domain, selector, and signing key name. |



---

|             |                                                                                  |
|-------------|----------------------------------------------------------------------------------|
| <b>Note</b> | If you do not enter search terms, the search engine returns all domain profiles. |
|-------------|----------------------------------------------------------------------------------|

---

## Signing System-Generated Messages

You can choose whether to sign system-generated messages with a DKIM signature. The appliance will sign the following messages:

- Cisco IronPort Spam Quarantine notifications
- Content filter-generated notifications
- Configuration messages
- Support requests



**Procedure**

- 
- Step 1** Choose **Mail Policies > Signing Profiles**.
- Step 2** In the **DKIM Signing of System Generated Messages** section, click **Edit Settings**.
- Step 3** Select **On**.
- Step 4** Submit and commit your changes.
- 

## Domain Keys and Logging

Lines such as the following are added to the mail logs upon DomainKeys signing:

```
Tue Aug 28 15:29:30 2007 Info: MID 371 DomainKeys: signing with dk-profile - matches
user123@example.com
Tue Aug 28 15:34:15 2007 Info: MID 373 DomainKeys: cannot sign - no profile matches
user12@example.com
```

Lines such as these are added to the mail logs upon DKIM signing:

```
Tue Aug 28 15:29:54 2007 Info: MID 372 DKIM: signing with dkim-profile - matches
user@example.com
Tue Aug 28 15:34:15 2007 Info: MID 373 DKIM: cannot sign - no profile matches
user2@example.com
```

## How to Verify Incoming Messages Using DKIM

**Table 20-1**      *How to Verify Incoming Messages Using DKIM*

|               | Do This                                                                                        | More Info                                                                                |
|---------------|------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Create a profile for verifying messages using DKIM.                                            | <a href="#">Creating a DKIM Verification Profile, page 20-18.</a>                        |
| <b>Step 2</b> | (Optional) Create a custom mail flow policy to use for verifying incoming messages using DKIM. | <a href="#">Defining Rules for Incoming Messages Using a Mail Flow Policy, page 7-15</a> |
| <b>Step 3</b> | Configure your mail flow policies to verify incoming messages using DKIM.                      | <a href="#">Configuring DKIM Verification on the Mail Flow Policy, page 20-20</a>        |
| <b>Step 4</b> | Define the action that the Email Security appliance takes on verified messages.                | <a href="#">Configuring an Action for DKIM Verified Mail, page 20-21</a>                 |
| <b>Step 5</b> | Associate the action with groups of specific senders or recipients.                            | <a href="#">Configuring Mail Policies, page 10-7</a>                                     |

**Related Topics**

- [DKIM Verification Checks Performed by AsyncOS, page 20-17](#)
- [Managing DKIM Verification Profiles, page 20-17](#)

- [Configuring DKIM Verification on the Mail Flow Policy, page 20-20](#)
- [Configuring an Action for DKIM Verified Mail, page 20-21](#)

## DKIM Verification Checks Performed by AsyncOS

When you configure an AsyncOS appliance for DKIM verification, the following checks are performed:

### Procedure

- 
- Step 1** AsyncOS checks for the DKIM-Signature field in incoming mail, the syntax of the signature header, valid tag values, and required tags. If the signature fails any of these checks, AsyncOS returns a *permfail*.
- Step 2** After the signature check is performed, the public key is retrieved from the public DNS record, and the TXT record is validated. If errors are encountered during this process, AsyncOS returns a *permfail*. A *tempfail* occurs if the DNS query for the public key fails to get a response.
- Step 3** After retrieving the public key, AsyncOS checks the hashed values and verifies the signature. If any failures occur during this step, AsyncOS returns a *permfail*.
- Step 4** If the checks all pass, AsyncOS returns a *pass*.



#### Note

When the message body is greater than the specified length, AsyncOS returns the following verdict:

```
dkim = pass (partially verified [x bytes])
```

where X represents the number of bytes verified.

The final verification result is entered as an *Authentication-Results* header. For example, you might get a header that looks like one of the following:

```
Authentication-Results: example1.com
```

```
header.from=From:user123@example.com; dkim=pass (signature verified)
```

```
Authentication-Results: example1.com
```

```
header.from=From:user123@example.com; dkim=pass (partially verified [1000 bytes])
```

```
Authentication-Results: example1.com
```

```
header.from=From:user123@example.com; dkim=permfail (body hash did not verify)
```



#### Note

Current DKIM verification stops at the first valid signature. It is not possible to verify using the last signature encountered. This functionality may be available in a later release.

## Managing DKIM Verification Profiles

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

- A name for the verification profile.
- The smallest and largest acceptable public key size. The default key sizes are 512 and 2048, respectively.
- The maximum number of signatures in the message to verify. If a message has more signatures than the maximum amount you defined, the appliance skips verification of the remaining signatures and continues to process the message. The default is 5 signatures.
- The maximum allowed difference in time (in seconds) between the sender's system time and verifier's. For example, if the message signature expires at 05:00:00 and the verifier's system time is 05:00:30, the message signature is still valid if the allowed difference in time is 60 seconds but it is invalid if the allowed difference is 10 seconds. The default is 60 seconds.
- An option whether to use a body length parameter.
- The SMTP action to take in case of a temporary failure.
- The SMTP action to take in case of a permanent failure.

You can search through all of your existing verification profiles by the profile name.

You can export your DKIM verification profiles as a text file in your appliance's configure directory. When you export the verification profiles, all of the profiles existing on the appliance are put into a single text file. See [Exporting DKIM Verification Profiles, page 20-19](#) for more information.

You can import DKIM verification profiles that you previously exported. Importing DKIM verification profiles causes all of the current DKIM verification profiles on the machine to be replaced. See [Importing DKIM Verification Profiles, page 20-19](#) for more information.

#### Related Topics

- [Creating a DKIM Verification Profile, page 20-18](#)
- [Exporting DKIM Verification Profiles, page 20-19](#)
- [Importing DKIM Verification Profiles, page 20-19](#)
- [Deleting DKIM Verification Profiles, page 20-19](#)
- [Searching DKIM Verification Profiles, page 20-20](#)

## Creating a DKIM Verification Profile

### Procedure

- 
- |               |                                                                                                                                    |
|---------------|------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Click <b>Mail Policies &gt; Verification Profiles</b> .                                                                            |
| <b>Step 2</b> | Click <b>Add Profile</b> .                                                                                                         |
| <b>Step 3</b> | Enter the name of the profile.                                                                                                     |
| <b>Step 4</b> | Select the minimum key size you want the appliance to accept for signing keys.                                                     |
| <b>Step 5</b> | Select the maximum key size you want the appliance to accept for signing keys.                                                     |
| <b>Step 6</b> | Select the maximum number of signatures to verify in a single message. The default is five signatures.                             |
| <b>Step 7</b> | Select the number of seconds before the key query times out. The default is 10 seconds.                                            |
| <b>Step 8</b> | Select maximum allowed difference in time (in seconds) between the sender's system time and verifier's. The default is 60 seconds. |
| <b>Step 9</b> | Select whether to use the body-length parameter in the signature to verify the message.                                            |

- Step 10** Select whether the Email Security appliance accepts or rejects the message if there is a temporary failure when verifying its signature. If you want the appliance to reject the message, you can choose to have it send the default 451 SMTP response code or another SMTP response code and text.
- Step 11** Select whether the Email Security appliance accepts or rejects the message if there is a permanent failure when verifying its signature. If you want the appliance to reject the message, you can choose to have it send the default 451 SMTP response code or another SMTP response code and text.
- Step 12** Submit your changes.
- The new profile appears in the DKIM Verification Profiles table.
- Step 13** Commit your changes.
- Step 14** At this point you should enable DKIM verification on an incoming mail flow policy and select the verification profile you want to use.
- 

## Exporting DKIM Verification Profiles

All DKIM verification profiles on the appliance are exported as a single text file and saved in the `configuration` directory on the appliance.

### Procedure

- 
- Step 1** Choose **Mail Policies > Verification Profiles**.
- Step 2** Click **Export Profiles**.
- Step 3** Enter a name for the file and click **Submit**.
- 

## Importing DKIM Verification Profiles

### Procedure

- 
- Step 1** Choose **Mail Policies > Verification Profiles**.
- Step 2** Click **Import Profiles**.
- Step 3** Select the file that contains the DKIM verification profiles.
- Step 4** Click **Submit**. You are warned that importing will replace all existing DKIM verification profiles.
- Step 5** Click **Import**.
- 

## Deleting DKIM Verification Profiles

### Related Topics

- [Removing Selected DKIM Verification Profiles, page 20-20](#)
- [Removing All DKIM Verification Profiles, page 20-20](#)

## Removing Selected DKIM Verification Profiles

### Procedure

---

- Step 1** Choose **Mail Policies > Verification Profiles**.
  - Step 2** Mark the checkbox to the right of each DKIM verification profile you want to delete.
  - Step 3** Click **Delete**.
  - Step 4** Confirm the deletion.
- 

## Removing All DKIM Verification Profiles

### Procedure

---

- Step 1** Choose **Mail Policies > Verification Profiles**.
  - Step 2** Click **Clear All Profiles**.
  - Step 3** Confirm the deletion.
- 

## Searching DKIM Verification Profiles

To search all DKIM verification profiles for a specific term in the profile name:

### Procedure

---

- Step 1** Choose **Mail Policies > Verification Profiles**.
  - Step 2** In the **Search DKIM Verification Profiles** section, specify the search term.
  - Step 3** Click **Find Profiles**.
    - The search scans the profile name for each DKIM verification profile.
    - If you do not enter search terms, the search engine returns all DKIM verification profiles.
- 

## Configuring DKIM Verification on the Mail Flow Policy

DKIM verification is enabled on mail flow policies for incoming email.

### Procedure

---

- Step 1** Choose **Mail Policies > Mail Flow Policies**.
- Step 2** Click the incoming mail policy for the listener where you want to perform verification.
- Step 3** In the Security Features section of the mail flow policy, enable DKIM Verification by selecting **On**.
- Step 4** Select the DKIM verification profile that you want to use for the policy.

**Step 5** Commit your changes.

---

#### Related Topics

- [DKIM Verification and Logging, page 20-21](#)

## DKIM Verification and Logging

Lines such as the following are added to the mail logs upon DKIM verification:

```
mail.current:Mon Aug 6 13:35:38 2007 Info: MID 17 DKIM: no signature
```

```
mail.current:Mon Aug 6 15:00:37 2007 Info: MID 18 DKIM: verified pass
```

## Configuring an Action for DKIM Verified Mail

When you verify DKIM mail, an *Authentication-Results* header is added to the mail, but the mail is accepted regardless of the authentication result. To configure actions based on these authentication results, you can create a content filter to perform actions on the DKIM-verified mail. For example, if DKIM verification fails, you may want configure the mail to be delivered, bounced, dropped, or sent to a quarantine. To do this, you must configure an action using a content filter.

#### Procedure

---

- Step 1** Choose **Mail Policies > Incoming Filters**.
- Step 2** Click **Add Filter**.
- Step 3** In the Conditions section, click **Add Condition**.
- Step 4** Select **DKIM Authentication** from the list of conditions.
- Step 5** Choose a DKIM condition. Select one of the following options:
- **Pass**. The message passed the authentication tests.
  - **Neutral**. Authentication was not performed.
  - **Temperror**. A recoverable error occurred.
  - **Permerror**. An unrecoverable error occurred.
  - **Hardfail**. The authentication tests failed.
  - **None**. The message was not signed.
- Step 6** Select an action to associate with the condition. For example, if the DKIM verification fails, you may want to notify the recipient and bounce the message. Or, if DKIM verification passes, you may want to deliver the message immediately without further processing.
- Step 7** Submit the new content filter.
- Step 8** Enable the content filter on the appropriate incoming mail policy.

**Step 9** Commit your changes.

---

## Overview of SPF and SIDF Verification

AsyncOS supports Sender Policy Framework (SPF) and Sender ID Framework (SIDF) verification. SPF and SIDF are methods for verifying authenticity of email based on DNS records. SPF and SIDF allow the owner of an Internet domain to use a special format of DNS TXT records to specify which machines are authorized to transmit email for that domain. Compliant mail receivers then use the published SPF records to test the authorization of the sending Mail Transfer Agent's identity during a mail transaction.

When you use SPF/SIDF authentication, the senders publish SPF records specifying which hosts are permitted to use their names, and compliant mail receivers use the published SPF records to test the authorization of the sending Mail Transfer Agent's identity during a mail transaction.



### Note

Because SPF checks require parsing and evaluation, AsyncOS performance may be impacted. In addition, be aware that SPF checks increase the load on your DNS infrastructure.

When you work with SPF and SIDF, note that SIDF is similar to SPF, but it has some differences. To get a full description of the differences between SIDF and SPF, see RFC 4406. For the purposes of this documentation, the two terms are discussed together except in the cases where only one type of verification applies.



### Note

AsyncOS does not support SPF for incoming relays.

### Related Topics

- [A Note About Valid SPF Records, page 20-22](#)

## A Note About Valid SPF Records

To use SPF and SIDF with a appliance, publish the SPF record according to the RFCs 4406 and 4408. Review RFC 4407 for a definition of how the PRA identity is determined. You may also want to refer to the following website to view common mistakes made when creating SPF and SIDF records:

[http://www.openspf.org/FAQ/Common\\_mistakes](http://www.openspf.org/FAQ/Common_mistakes)

### Related Topics

- [Valid SPF Records, page 20-22](#)
- [Valid SIDF Records, page 20-23](#)
- [Testing Your SPF Records, page 20-23](#)

## Valid SPF Records

To pass the SPF HELO check, ensure that you include a “v=spf1 a –all” SPF record for each sending MTA (separate from the domain). If you do not include this record, the HELO check will likely result in a None verdict for the HELO identity. If you notice that SPF senders to your domain return a high number of None verdicts, these senders may not have included a “v=spf1 a –all” SPF record for each sending MTA.

## Valid SIDF Records

To support the SIDF framework, you need to publish both “v=spf1” and “spf2.0” records. For example, your DNS record may look like the following example:

```
example.com. TXT "v=spf1 +mx a:colo.example.com/28 -all"

smtp-out.example.com TXT "v=spf1 a -all"

example.com. TXT "spf2.0/mfrom,pra +mx a:colo.example.com/28 -all"
```

SIDF does not verify the HELO identity, so in this case, you do not need to publish SPF v2.0 records for each sending MTA.



### Note

If you choose not to support SIDF, publish an “spf2.0/pra ~all” record.

## Testing Your SPF Records

In addition to reviewing the RFCs, it is a good idea to test your SPF records before you implement SPF verification on an Email Security appliance. There are several testing tools available on the [openspf.org](http://www.openspf.org) website:

<http://www.openspf.org/Tools>

You can use the following tool to determine why an email failed an SPF record check:

<http://www.openspf.org/Why>

In addition, you can enable SPF on a test listener and use Cisco’s `trace` CLI command (or perform trace from the GUI) to view the SPF results. Using trace, you can easily test different sending IPs.

# How to Verify Incoming Messages Using SPF/SDIF

**Table 20-2**      *How to Verify Incoming Messages Using SPF/SDIF*

|               | Do This                                                                                            | More Info                                                                     |
|---------------|----------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------|
| <b>Step 1</b> | (Optional) Create a custom mail flow policy to use for verifying incoming messages using SPF/SDIF. | <a href="#">Defining Rules for Incoming Messages Using Policy, page 7-15</a>  |
| <b>Step 2</b> | Configure your mail flow policies to verify incoming messages using SPF/SDIF.                      | <a href="#">Enabling SPF and SIDF, page 20-24</a>                             |
| <b>Step 3</b> | Define the action that the Email Security appliance takes on verified messages.                    | <a href="#">Determining the Action to Take for SPF/SIDF Mail, page 20-31.</a> |
| <b>Step 4</b> | Associate the action with groups of specific senders or recipients.                                | <a href="#">Configuring Mail Policies, page 10-7</a>                          |
| <b>Step 5</b> | (Optional) Test the results of message verification.                                               | <a href="#">Testing the SPF/SIDF Results, page 20-34</a>                      |



**Caution**

Although Cisco strongly endorses email authentication globally, at this point in the industry's adoption, Cisco suggests a cautious disposition for SPF/SIDF authentication failures. Until more organizations gain greater control of their authorized mail sending infrastructure, Cisco urges customers to avoid bouncing emails and instead quarantine emails that fail SPF/SIDF verification.

**Note**

The AsyncOS command line interface (CLI) provides more control settings for SPF level than the web interface. Based on the SPF verdict, the appliance can accept or reject a message, in SMTP conversation, on a per listener basis. You can modify the SPF settings when editing the default settings for a listener's Host Access Table using the `listenerconfig` command. See the [Enabling SPF and SIDF via the CLI, page 20-25](#) for more information on the settings.

## Enabling SPF and SIDF

To use SPF/SIDF, you must enable SPF/SIDF for a mail flow policy on an incoming listener. You can enable SPF/SIDF on the listener from the default mail flow policy, or you can enable it for particular incoming mail flow policies.

### Procedure

- Step 1** Choose **Mail Policies > Mail Flow Policy**.
- Step 2** Click **Default Policy Parameters**.
- Step 3** In the default policy parameters, view the Security Features section.
- Step 4** In the **SPF/SIDF Verification** section, click **On**.
- Step 5** Set the level of conformance (the default is SIDF-compatible). This option allows you to determine which standard of SPF or SIDF verification to use. In addition to SIDF conformance, you can choose SIDF-compatible, which combines SPF and SIDF.

**Table 20-3**      **SPF/SIDF Conformance Levels**

| Conformance Level | Description                                                                                                                                                                                                              |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SPF               | The SPF/SIDF verification behaves according to RFC4408.<br>- No purported responsible address (PRA) identity verification takes place.<br><b>NOTE:</b> Select this conformance option to test against the HELO identity. |

**Table 20-3**      **SPF/SIDF Conformance Levels**

| Conformance Level | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SIDF              | <p>The SPF/SIDF verification behaves according to RFC4406.</p> <ul style="list-style-type: none"> <li>-The PRA Identity is determined with full conformance to the standard.</li> <li>- SPF v1.0 records are treated as spf2.0/mfrom,pra.</li> <li>- For a nonexistent domain or a malformed identity, a verdict of Fail is returned.</li> </ul>                                                                                                                    |
| SIDF Compatible   | <p>The SPF/SIDF verification behaves according to RFC4406 <i>except for</i> the following differences:</p> <ul style="list-style-type: none"> <li>- SPF v1.0 records are treated as spf2.0/mfrom.</li> <li>- For a nonexistent domain or a malformed identity, a verdict of None is returned.</li> </ul> <p><b>NOTE:</b> This conformance option was introduced at the request of the OpenSPF community (<a href="http://www.openspf.org">www.openspf.org</a>).</p> |



**Note** More settings are available via the CLI. See [Enabling SPF and SIDF via the CLI, page 20-25](#) for more information.

- Step 6** If you choose a conformance level of SIDF-compatible, configure whether the verification downgrades a Pass result of the PRA identity to None if there are Resent-Sender: or Resent-From: headers present in the message. You might choose this option for security purposes.
- Step 7** If you choose a conformance level of SPF, configure whether to perform a test against the HELO identity. You might use this option to improve performance by disabling the HELO check. This can be useful because the `spf-passed` filter rule checks the PRA or the MAIL FROM Identities first. The appliance only performs the HELO check for the SPF conformance level.

#### Related Topics

- [Enabling SPF and SIDF via the CLI, page 20-25](#)
- [The Received-SPF Header, page 20-30](#)

## Enabling SPF and SIDF via the CLI

The AsyncOS CLI supports more control settings for each SPF/SIDF conformance level. When configuring the default settings for a listener's Host Access Table, you can choose the listener's SPF/SIDF conformance level and the SMTP actions (ACCEPT or REJECT) that the appliance performs, based on the SPF/SIDF verification results. You can also define the SMTP response that the appliance sends when it rejects a message.

Depending on the conformance level, the appliance performs a check against the HELO identity, MAIL FROM identity, or PRA identity. You can specify whether the appliance proceeds with the session (ACCEPT) or terminates the session (REJECT) for each of the following SPF/SIDF verification results for each identity check:

- **None.** No verification can be performed due to the lack of information.

- **Neutral.** The domain owner does not assert whether the client is authorized to use the given identity.
- **SoftFail.** The domain owner believes the host is not authorized to use the given identity but is not willing to make a definitive statement.
- **Fail.** The client is not authorized to send mail with the given identity.
- **TempError.** A transient error occurred during verification.
- **PermError.** A permanent error occurred during verification.

The appliance accepts the message for a Pass result unless you configure the SIDF Compatible conformance level to downgrade a Pass result of the PRA identity to None if there are Resent-Sender: or Resent-From: headers present in the message. The appliance then takes the SMTP action specified for when the PRA check returns None.

If you choose not to define the SMTP actions for an identity check, the appliance automatically accepts all verification results, including Fail.

The appliance terminates the session if the identity verification result matches a REJECT action for any of the enabled identity checks. For example, an administrator configures a listener to accept messages based on all HELO identity check results, including Fail, but also configures it to reject messages for a Fail result from the MAIL FROM identity check. If a message fails the HELO identity check, the session proceeds because the appliance accepts that result. If the message then fails the MAIL FROM identity check, the listener terminates the session and then returns the SMTP response for the REJECT action.

The SMTP response is a code number and message that the appliance returns when it rejects a message based on the SPF/SIDF verification result. The TempError result returns a different SMTP response from the other verification results. For TempError, the default response code is 451 and the default message text is #4.4.3 Temporary error occurred during SPF verification. For all other verification results, the default response code is 550 and the default message text is #5.7.1 SPF unauthorized mail is prohibited. You can specify your own response code and message text for TempError and the other verification results.

Optionally, you can configure the appliance to return a third-party response from the SPF publisher domain if the REJECT action is taken for Neutral, SoftFail, or Fail verification result. By default, the appliance returns the following response:

```
550-#5.7.1 SPF unauthorized mail is prohibited.

550-The domain example.com explains:

550 <Response text from SPF domain publisher>
```

To enable these SPF/SIDF settings, use the `listenerconfig -> edit` subcommand and select a listener. Then use the `hostaccess -> default` subcommand to edit the Host Access Table's default settings. Answer `yes` to the following prompts to configure the SPF controls:

```
Would you like to change SPF/SIDF settings? [N]> yes

Would you like to perform SPF/SIDF Verification? [Y]> yes
```

The following SPF control settings are available for the Host Access Table:

**Table 20-4** *SPF Control Settings via the CLI*

| Conformance Level | Available SPF Control Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SPF Only          | <ul style="list-style-type: none"> <li>• whether to perform HELO identity check</li> <li>• SMTP actions taken based on the results of the following identity checks: <ul style="list-style-type: none"> <li>– HELO identity (if enabled)</li> <li>– MAIL FROM Identity</li> </ul> </li> <li>• SMTP response code and text returned for the REJECT action</li> <li>• verification time out (in seconds)</li> </ul>                                                                                                                                                                                            |
| SIDF Compatible   | <ul style="list-style-type: none"> <li>• whether to perform a HELO identity check</li> <li>• whether the verification downgrades a Pass result of the PRA identity to None if the Resent-Sender: or Resent-From: headers are present in the message</li> <li>• SMTP actions taken based on the results of the following identity checks: <ul style="list-style-type: none"> <li>– HELO identity (if enabled)</li> <li>– MAIL FROM Identity</li> <li>– PRA Identity</li> </ul> </li> <li>• SMTP response code and text returned for the REJECT action</li> <li>• verification timeout (in seconds)</li> </ul> |
| SIDF Strict       | <ul style="list-style-type: none"> <li>• SMTP actions taken based on the results of the following identity checks: <ul style="list-style-type: none"> <li>– MAIL FROM Identity</li> <li>– PRA Identity</li> </ul> </li> <li>• SMTP response code and text returned in case of SPF REJECT action</li> <li>• verification timeout (in seconds)</li> </ul>                                                                                                                                                                                                                                                      |

The following example shows a user configuring the SPF/SIDF verification using the SPF Only conformance level. The appliance performs the HELO identity check and accepts the None and Neutral verification results and rejects the others. The CLI prompts for the SMTP actions are the same for all identity types. The user does not define the SMTP actions for the MAIL FROM identity. The appliance automatically accepts all verification results for the identity. The appliance uses the default reject code and text for all REJECT results.

```
Would you like to change SPF/SIDF settings? [N]> yes
```

```
Would you like to perform SPF/SIDF Verification? [N]> yes
```

What Conformance Level would you like to use?

1. SPF only
2. SIDF compatible
3. SIDF strict

[2]> **1**

Would you like to have the HELO check performed? [Y]> **y**

Would you like to change SMTP actions taken as result of the SPF verification? [N]> **y**

Would you like to change SMTP actions taken for the HELO identity? [N]> **y**

What SMTP action should be taken if HELO check returns None?

1. Accept
2. Reject

[1]> **1**

What SMTP action should be taken if HELO check returns Neutral?

1. Accept
2. Reject

[1]> **1**

What SMTP action should be taken if HELO check returns SoftFail?

1. Accept
2. Reject

[1]> **2**

What SMTP action should be taken if HELO check returns Fail?

1. Accept

2. Reject

[1]> **2**

What SMTP action should be taken if HELO check returns TempError?

1. Accept

2. Reject

[1]> **2**

What SMTP action should be taken if HELO check returns PermError?

1. Accept

2. Reject

[1]> **2**

Would you like to change SMTP actions taken for the MAIL FROM identity? [N]> **n**

Would you like to change SMTP response settings for the REJECT action? [N]> **n**

Verification timeout (seconds)

[40]>

The following shows how the SPF/SIDF settings are displayed for the listener's Default Policy Parameters.

SPF/SIDF Verification Enabled: Yes

Conformance Level: SPF only

Do HELO test: Yes

SMTP actions:

For HELO Identity:

None, Neutral: Accept

```

SoftFail, Fail, TempError, PermError: Reject

For MAIL FROM Identity: Accept

SMTP Response Settings:

Reject code: 550

Reject text: #5.7.1 SPF unauthorized mail is prohibited.

Get reject response text from publisher: Yes

Defer code: 451

Defer text: #4.4.3 Temporary error occurred during SPF verification.

Verification timeout: 40

```

See the *Cisco AsyncOS CLI Reference Guide* for more information on the `listenerconfig` command.

## The Received-SPF Header

When you configure AsyncOS for SPF/SIDF verification, it places an SPF/SIDF verification header (Received-SPF) in the email. The Received-SPF header contains the following information:

- **verification result** - the SPF verification result (see [Verification Results, page 20-31](#)).
- **identity** - the identity that SPF verification checked: HELO, MAIL FROM, or PRA.
- **receiver** - the verifying host name (which performs the check).
- **client IP address** - the IP address of the SMTP client.
- **ENVELOPE FROM** - the envelope sender mailbox. (Note that this may be different from the MAIL FROM identity, as the MAIL FROM identity cannot be empty.)
- **x-sender** - the value of the HELO, MAIL FROM, or PRA identity.
- **x-conformance** - the level of conformance (see [Table 20-3SPF/SIDF Conformance Levels, page 20-24](#)) and whether a downgrade of the PRA check was performed.

The following example shows a header added for a message that passed the SPF/SIDF check:

```

Received-SPF: Pass identity=pra; receiver=box.example.com;

client-ip=1.2.3.4; envelope-from="alice@fooo.com";

x-sender="alice@company.com"; x-conformance=sidf_compatible

```



### Note

The `spf-status` and `spf-passed` filter rules use the received-SPF header to determine the status of the SPF/SIDF verification.

# Determining the Action to Take for SPF/SIDF Verified Mail

When you receive SPF/SIDF verified mail, you may want to take different actions depending on the results of the SPF/SIDF verification. You can use the following message and content filter rules to determine the status of SPF/SIDF verified mail and perform actions on the messages based on the verification results:

- `spf-status`. This filter rule determines actions based on the SPF/SIDF status. You can enter a different action for each valid SPF/SIDF return value.
- `spf-passed`. This filter rule generalizes the SPF/SIDF results as a Boolean value.



**Note** The `spf-passed` filter rule is only available in message filters.

You can use the `spf-status` rule when you want to address more granular results, and use the `spf-passed` rule when you want to create a simple Boolean.

## Related Topics

- [Verification Results, page 20-31](#)
- [Using the `spf-status` Filter Rule in the CLI, page 20-32](#)
- [spf-status Content Filter Rule in the GUI, page 20-33](#)
- [Using the `spf-passed` Filter Rule, page 20-33](#)

## Verification Results

If you use the `spf-status` filter rule, you can check against the SPF/SIDF verification results using the following syntax:

```
if (spf-status == "Pass")
```

If you want a single condition to check against multiple status verdicts, you can use the following syntax:

```
if (spf-status == "PermError, TempError")
```

You can also check the verification results against the HELO, MAIL FROM, and PRA identities using the following syntax:

```
if (spf-status("pra") == "Fail")
```



**Note**

You can only use the `spf-status` message filter rule to check results against HELO, MAIL FROM, and PRA identities. You cannot use the `spf-status` content filter rule to check against identities. The `spf-status` content filter checks only the PRA identity.

You can receive any of the following verification results:

- None - no verification can be performed due to the lack of information.



- Pass - the client is authorized to send mail with the given identity.
- Neutral - the domain owner does not assert whether the client is authorized to use the given identity.
- SoftFail - the domain owner believes the host is not authorized to use the given identity but is not willing to make a definitive statement.
- Fail - the client is not authorized to send mail with the given identity.
- TempError - a transient error occurred during verification.
- PermError - a permanent error occurred during verification.

## Using the spf-status Filter Rule in the CLI

The following example shows the `spf-status` message filter in use:

```
skip-spam-check-for-verified-senders:

 if (sendergroup == "TRUSTED" and spf-status == "Pass"){

 skip-spamcheck();

 }

quarantine-spf-failed-mail:

 if (spf-status("pra") == "Fail") {

 if (spf-status("mailfrom") == "Fail"){

 # completely malicious mail

 quarantine("Policy");

 } else {

 if(spf-status("mailfrom") == "SoftFail") {

 # malicious mail, but tempting

 quarantine("Policy");

 }

 }

 } else {

 if(spf-status("pra") == "SoftFail"){

 if (spf-status("mailfrom") == "Fail"

 or spf-status("mailfrom") == "SoftFail"){

 # malicious mail, but tempting
```

```

 quarantine("Policy");
 }
}

stamp-mail-with-spf-verification-error:

 if (spf-status("pra") == "PermError, TempError"

 or spf-status("mailfrom") == "PermError, TempError"

 or spf-status("helo") == "PermError, TempError"){

 # permanent error - stamp message subject

 strip-header("Subject");

 insert-header("Subject", "[POTENTIAL PHISHING] $Subject"); }

```

## spf-status Content Filter Rule in the GUI

You can also enable the `spf-status` rule from the content filters in the GUI. However, you cannot check results against HELO, MAIL FROM, and PRA identities when using the `spf-status` content filter rule.

To add the `spf-status` content filter rule from the GUI, click **Mail Policies > Incoming Content Filters**. Then add the SPF Verification filter rule from the Add Condition dialog box. Specify one or more verification results for the condition.

After you add the SPF Verification condition, specify an action to perform based on the SPF status. For example, if the SPF status is SoftFail, you might quarantine the message.

## Using the spf-passed Filter Rule

The `spf-passed` rule shows the results of SPF verification as a Boolean value. The following example shows an `spf-passed` rule used to quarantine emails that are not marked as `spf-passed`:

```

quarantine-spf-unauthorized-mail:

 if (not spf-passed) {

 quarantine("Policy");

 }

```

**Note**

Unlike the `spf-status` rule, the `spf-passed` rule reduces the SPF/SIDF verification values to a simple Boolean. The following verification results are treated as not passed in the `spf-passed` rule: None, Neutral, Softfail, TempError, PermError, and Fail. To perform actions on messages based on more granular results, use the `spf-status` rule.

## Testing the SPF/SIDF Results

Test the results of SPF/SIDF verification and use these results to determine how to treat SPF/SIDF failures because different organizations implement SPF/SIDF in different ways. Use a combination of content filters, message filters, and the Email Security Monitor - Content Filters report to test the results of the SPF/SIDF verification.

Your degree of dependence on SPF/SIDF verification determines the level of granularity at which you test SPF/SIDF results.

**Related Topics**

- [Basic Granularity Test of SPF/SIDF Results, page 20-34](#)
- [Greater Granularity Test of SPF/SIDF Results, page 20-34](#)

## Basic Granularity Test of SPF/SIDF Results

To get a basic measure of the SPF/SIDF verification results for incoming mail, you can use content filters and the Email Security Monitor - Content Filters page. This test provides a view of the number of messages received for each type of SPF/SIDF verification result.

**Procedure**

- |               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Enable SPF/SIDF verification for a mail flow policy on an incoming listener, and use a content filter to configure an action to take. For information on enabling SPF/SIDF, see <a href="#">Enabling SPF and SIDF, page 20-24</a> .                                                                                                                                                                                                                                                      |
| <b>Step 2</b> | Create an <code>spf-status</code> content filter for each type of SPF/SIDF verification. Use a naming convention to indicate the type of verification. For example, use “SPF-Passed” for messages that pass SPF/SIDF verification, or “SPF-TempErr” for messages that weren’t passed due to a transient error during verification. For information about creating an <code>spf-status</code> content filter, see <a href="#">spf-status Content Filter Rule in the GUI, page 20-33</a> . |
| <b>Step 3</b> | After you have processed a number of SPF/SIDF verified messages, click Monitor > Content Filters to see how many messages triggered each of the SPF/SIDF verified content filters.                                                                                                                                                                                                                                                                                                       |

## Greater Granularity Test of SPF/SIDF Results

For more comprehensive information about SPF/SIDF verification results, only enable SPF/SIDF verification for specific groups of senders, and review the results for those specific senders. Then, create a mail policy for that particular group and enable SPF/SIDF verification on the mail policy. Create

content filters and review the Content Filters report as explained in [Basic Granularity Test of SPF/SIDF Results, page 20-34](#). If you find that the verification is effective, then you can use SPF/SIDF verification as a basis for deciding whether to drop or bounce emails for this specified group of senders.

#### Procedure

- 
- |               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Create a mail flow policy for SPF/SIDF verification. Enable SPF/SIDF verification for the mail flow policy on an incoming listener. For information about enabling SPF/SIDF, see <a href="#">Enabling SPF and SIDF, page 20-24</a> .                                                                                                                                                                                                                                                     |
| <b>Step 2</b> | Create a sender group for SPF/SIDF verification and use a naming convention to indicate SPF/SIDF verification. For information about creating sender groups, see the “Configuring the Gateway to Receive Mail” chapter.                                                                                                                                                                                                                                                                  |
| <b>Step 3</b> | Create an <code>spf-status</code> content filter for each type of SPF/SIDF verification. Use a naming convention to indicate the type of verification. For example, use “SPF-Passed” for messages that pass SPF/SIDF verification, or “SPF-TempErr” for messages that weren’t passed due to a transient error during verification. For information about creating an <code>spf-status</code> content filter, see <a href="#">spf-status Content Filter Rule in the GUI, page 20-33</a> . |
| <b>Step 4</b> | After you process a number of SPF/SIDF-verified messages, click Monitor > Content Filters to see how many messages triggered each of the SPF/SIDF-verified content filters.                                                                                                                                                                                                                                                                                                              |
- 

## DMARC Verification

Domain-based Message Authentication, Reporting and Conformance (DMARC) is a technical specification created to reduce the potential for email-based abuse. DMARC standardizes how email receivers perform email authentication using SPF and DKIM mechanisms. To pass DMARC verification, an email must pass at least one of these authentication mechanisms, and the Authentication Identifiers must comply with RFC 5322.

AsyncOS for Email allows you to:

- Verify incoming emails using DMARC.
- Define profiles to override (accept, quarantine, or reject) domain owners’ policies.
- Send feedback reports to domain owners, which helps to strengthen their authentication deployments.
- Send delivery error reports to the domain owners if the DMARC aggregate report size exceeds 10 MB or the size specified in the RUA tag of the DMARC record.

AsyncOS for Email can handle emails that are compliant with the DMARC specification as submitted to Internet Engineering Task Force (IETF) on March 31, 2013. For more information, see <http://tools.ietf.org/html/draft-kucherawy-dmarc-base-02>.

#### Related Topics

- [DMARC Verification Workflow in AsyncOS for Email, page 20-36](#)
- [How to Verify Incoming Messages Using DMARC, page 20-36](#)

## DMARC Verification Workflow in AsyncOS for Email

The following describes how AsyncOS for Email performs DMARC verification.

1. A listener configured on AsyncOS receives an SMTP connection.
2. AsyncOS performs SPF and DKIM verification on the message.
3. AsyncOS fetches the DMARC record for the sender's domain from the DNS.
  - If no record is found, AsyncOS skips the DMARC verification and continues processing.
  - If the DNS lookup fails, AsyncOS takes action based on the specified DMARC verification profile.
4. Depending on DKIM and SPF verification results, AsyncOS performs DMARC verification on the message.



**Note** If DKIM and SPF verification is enabled, DMARC verification reuses the DKIM and SPF verification results.

5. Depending on the DMARC verification result and the specified DMARC verification profile, AsyncOS accepts, quarantines, or rejects the message. If the message is not rejected due to DMARC verification failure, AsyncOS continues processing.
6. AsyncOS sends an appropriate SMTP response and continues processing.
7. If sending of aggregate reports is enabled, AsyncOS gathers DMARC verification data and includes it in the daily report sent to the domain owners. For more information about the DMARC aggregate feedback report, see [DMARC Aggregate Reports, page 20-42](#).



**Note** If the aggregate report size exceeds 10 MB or the size specified in the RUA tag of the DMARC record, AsyncOS sends delivery error reports to the domain owners.

## How to Verify Incoming Messages Using DMARC

**Table 20-5**      *How to Verify Incoming Messages Using DMARC*

|               | Do This                                                                                                             | More Information                                                                                                                 |
|---------------|---------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Create a new DMARC verification profile or modify the default DMARC verification profile to meet your requirements. | <a href="#">Create a DMARC Verification Profile, page 20-37</a><br><a href="#">Edit a DMARC Verification Profile, page 20-39</a> |
| <b>Step 2</b> | (Optional) Configure global DMARC settings to meet your requirements.                                               | <a href="#">Configure Global DMARC Settings, page 20-40</a>                                                                      |
| <b>Step 3</b> | Configure your mail flow policies to verify incoming messages using DMARC.                                          | <a href="#">Configuring DMARC Verification on the Mail Flow Policy, page 20-41</a>                                               |

**Table 20-5**      *How to Verify Incoming Messages Using DMARC*

|               | Do This                                                                                                                                                                                                     | More Information                                                                                                                                                                                                              |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 4</b> | (Optional) Configure a return address for DMARC feedback reports.                                                                                                                                           | <a href="#">Configure a Return Address for DMARC Feedback Reports, page 20-42</a>                                                                                                                                             |
| <b>Step 5</b> | (Optional) Review the following: <ul style="list-style-type: none"> <li>• DMARC Verification and Incoming Mail reports</li> <li>• Messages that failed DMARC verification using Message Tracking</li> </ul> | <ul style="list-style-type: none"> <li>• <a href="#">DMARC Verification Page, page 28-19</a></li> <li>• <a href="#">Incoming Mail Page, page 28-9</a></li> <li>• <a href="#">Searching for Messages, page 29-2</a></li> </ul> |

**Related Topics**

- [Managing DMARC Verification Profiles, page 20-37](#)
- [Configure Global DMARC Settings, page 20-40](#)
- [Configuring DMARC Verification on the Mail Flow Policy, page 20-41](#)
- [Configure a Return Address for DMARC Feedback Reports, page 20-42](#)
- [DMARC Aggregate Reports, page 20-42](#)

## Managing DMARC Verification Profiles

A DMARC verification profile is a list of parameters that the mail flow policies of the Email Security appliance use for verifying DMARC. For example, you may want to create a stringent profile that rejects all non-compliant messages from a particular domain and a less stringent profile that quarantines all non-compliant messages from another domain.

A DMARC verification profile consists of the following information:

- A name for the verification profile.
- Message action to take when the policy in the DMARC record is reject.
- Message action to take when the policy in the DMARC record is quarantine.
- Message action in case of a temporary failure.
- Message action in case of a permanent failure.

**Related Topics**

- [Create a DMARC Verification Profile, page 20-37](#)
- [Edit a DMARC Verification Profile, page 20-39](#)
- [Exporting DMARC Verification Profiles, page 20-39](#)
- [Importing DMARC Verification Profiles, page 20-39](#)
- [Deleting DMARC Verification Profiles, page 20-39](#)

### Create a DMARC Verification Profile

Use this procedure to create a new DMARC verification profile.

**Note**

By default, AsyncOS provides a default DMARC verification profile. If you do not want to create a new DMARC verification profile, you can use the default DMARC verification profile. The default DMARC verification profile is available on **Mail Policies > DMARC** page. For instructions to edit the default DMARC verification profile, see [Edit a DMARC Verification Profile, page 20-39](#).

**Procedure**

- 
- Step 1** Choose **Mail Policies > DMARC**.
- Step 2** Click **Add Profile**.
- Step 3** Enter the name of the profile.
- Step 4** Set the message action that AsyncOS takes when the policy in the DMARC record is reject. Choose one of the following:
- **No Action.** AsyncOS does not take any action on the messages that fail DMARC verification.
  - **Quarantine.** AsyncOS quarantines the messages that fail DMARC verification to a specified quarantine.
  - **Reject.** AsyncOS rejects all messages that fail DMARC verification and returns a specified SMTP code and response. The default values are, respectively: 550 and #5.7.1 DMARC unauthenticated mail is prohibited.
- Step 5** Set the message action that AsyncOS takes when the policy in the DMARC record is quarantine. Choose one of the following:
- **No Action.** AsyncOS does not take any action on the messages that fail DMARC verification.
  - **Quarantine.** AsyncOS quarantines the messages that fail DMARC verification to a specified quarantine.
- Step 6** Set the message action that AsyncOS takes on the messages that result in temporary failure during DMARC verification. Choose one of the following:
- **Accept.** AsyncOS accepts messages that result in temporary failure during DMARC verification.
  - **Reject.** AsyncOS rejects messages that result in temporary failure during DMARC verification and returns a specified SMTP code and response. The default values are, respectively: 451 and #4.7.1 Unable to perform DMARC verification.
- Step 7** Set the message action that AsyncOS takes on the messages that result in permanent failure during DMARC verification. Choose one of the following:
- **Accept.** AsyncOS accepts messages that result in permanent failure during DMARC verification.
  - **Reject.** AsyncOS rejects messages that result in permanent failure during DMARC verification, and returns a specified SMTP code and response. The default values are, respectively: 550 and #5.7.1 DMARC verification failed.
- Step 8** Submit and commit your changes.
-

## Edit a DMARC Verification Profile

### Procedure

- 
- Step 1** Choose **Mail Policies > DMARC**.
  - Step 2** Click the intended verification profile name.
  - Step 3** Edit the intended fields as described in [Create a DMARC Verification Profile, page 20-37](#).
  - Step 4** Submit and commit your changes.
- 

## Exporting DMARC Verification Profiles

You can export all DMARC verification profiles on your appliance to a single text file in the `configuration` directory.

### Procedure

- 
- Step 1** Choose **Mail Policies > DMARC**.
  - Step 2** Click **Export Profiles**.
  - Step 3** Enter a name for the file.
  - Step 4** Click **Submit**.
- 

## Importing DMARC Verification Profiles

### Procedure

- 
- Step 1** Choose **Mail Policies > DMARC**.
  - Step 2** Click **Import Profiles**.
  - Step 3** Choose the file that contains the DMARC verification profiles.
  - Step 4** Click **Submit**. You are warned that importing will replace all existing DMARC verification profiles.
  - Step 5** Click **Import**.
  - Step 6** Commit your changes.
- 

## Deleting DMARC Verification Profiles

### Procedure

- 
- Step 1** Choose **Mail Policies > DMARC**.
  - Step 2** Select the verification profiles that you want to delete.
  - Step 3** Click **Delete**.



**Step 4** Confirm the deletion.

## Configure Global DMARC Settings

### Procedure

- Step 1** Choose **Mail Policies > DMARC**.
- Step 2** Click **Edit Global Settings**.
- Step 3** Make changes to the settings defined in the following table.

**Table 20-6** *DMARC Global Settings*

| Global Setting                                | Description                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Specific senders bypass address list          | <p>Skip DMARC verification of messages from specific senders. Choose an address list from the drop-down list.</p> <p><b>Note</b> You can choose only the address lists that are created by choosing the Allow only full Email Addresses option. For more information, see <a href="#">Using a List of Sender Addresses for Incoming Connection Rules</a>, page 7-22.</p> |
| Bypass verification for messages with headers | <p>Skip DMARC verification of messages that contain specific headers. For example, use this option to skip DMARC verification of messages from mailing lists and trusted forwarders.</p> <p>Enter a header or multiple headers separated by commas.</p>                                                                                                                  |
| Schedule for report generation                | <p>The time when you want AsyncOS to generate DMARC aggregate reports. For example, you can choose non-peak hours for generating aggregate reports to avoid impact on mail flow.</p>                                                                                                                                                                                     |
| Entity generating reports                     | <p>The entity generating DMARC aggregate reports. This helps the domain owners who receive DMARC aggregate reports to identify the entity that generated the report.</p> <p>Enter a valid domain name.</p>                                                                                                                                                               |
| Additional contact information for reports    | <p>Additional contact information, for example, details of your organization's customer support, if the domain owners who receive DMARC aggregate reports want to contact the entity that generated the report.</p>                                                                                                                                                      |

**Table 20-6**      **DMARC Global Settings**

| Global Setting                        | Description                                                                                                                                                                                                       |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Send copy of all aggregate reports to | Send a copy of all DMARC aggregate reports to specific users, for example, internal users who perform analysis on the aggregate reports.<br><br>Enter an email address or multiple addresses separated by commas. |
| Error Reports                         | Send delivery error reports to the domain owners if the DMARC aggregate report size exceeds 10 MB or the size specified in the RUA tag of the DMARC record.<br><br>Check the check box.                           |

**Step 4**      Submit and commit your changes.

## Configuring DMARC Verification on the Mail Flow Policy

### Procedure

- Step 1**      Choose **Mail Policies > Mail Flow Policies**.
- Step 2**      Click the incoming mail policy for the listener where you want to perform verification.
- Step 3**      In the Security Features section of the mail flow policy, enable DMARC Verification by choosing **On**.
- Step 4**      Select the DMARC verification profile that you want to use for the policy.
- Step 5**      (Optional) Enable sending of DMARC aggregate feedback reports to email addresses in the RUA tag of DMARC-enabled domains from whom messages are received.  
  
Aggregate feedback reports are generated daily.
- Step 6**      Submit and commit your changes.

### Related Topics

[DMARC Verification Logs, page 20-41](#)

## DMARC Verification Logs

Log messages are added to the mail logs during the following stages of DMARC verification.

- DMARC verification attempted on a message
- DMARC verification is complete
- DMARC verification details including DKIM and SPF alignment results
- DMARC verification on a message is skipped
- DMARC record is fetched and parsed or DNS failures
- DMARC aggregate report delivery for a domain failed
- Error report generated for a domain

- Error report delivery for a domain succeeded
- Error report delivery for a domain failed

## Configure a Return Address for DMARC Feedback Reports

### Procedure

- 
- Step 1** Choose **System Administration > Return Addresses**.
- Step 2** Click **Edit Settings**.
- Step 3** Provide a return address for DMARC aggregate feedback reports.
- Step 4** Submit and commit your changes.
- 

## DMARC Aggregate Reports

DMARC relies on a feedback mechanism to enforce domain owner policies safely and in a scalable manner. This feedback mechanism helps the domain owners to strengthen their authentication deployments.

If you are using AsyncOS to perform DMARC verification and you have enabled sending of aggregate feedback reports in the mail flow policy, AsyncOS generates aggregate feedback reports daily, and sends it to the domain owners. These reports are in XML format and are archived into a GZip file.



### Note

---

All DMARC aggregate feedback reports that AsyncOS generates are DMARC compliant.

---

A DMARC aggregate feedback report contains the following sections:

- Metadata of the report sender such as email address and report ID number.
- Details of the published DMARC policy.
- Details of DMARC policy disposition such as source IP address and disposition summary.
- Domain identifiers
- DMARC verification results and authentication summary.

### Related Topics

- [Sample DMARC Aggregate Feedback Report, page 20-42](#)

## Sample DMARC Aggregate Feedback Report

```
<?xml version="1.0" encoding="UTF-8" ?>
<feedback>
 <version>1.0</version>
 <report_metadata>
 <org_name>cisco.com</org_name>
 <email>noreply-dmarc-support@cisco.com</email>
 <extra_contact_info>http://cisco.com/dmarc/support</extra_contact_info>
 <report_id>b1d925$4ecceab=0694614b826605cd@cisco.com</report_id>
 <date_range>
 <begin>1335571200</begin>
 <end>1335657599</end>
```

```

 </date_range>
 </report_metadata>
 <policy_published>
 <domain>example.com</domain>
 <adkim>r</adkim>
 <aspf>r</aspf>
 <p>none</p>
 <sp>none</sp>
 <pct>100</pct>
 </policy_published>
 <record>
 <row>
 <source_ip>1.1.1.1</source_ip>
 <count>2</count>
 <policy_evaluated>
 <disposition>none</disposition>
 <dkim>fail</dkim>
 <spf>pass</spf>
 </policy_evaluated>
 </row>
 <identifiers>
 <envelope_from>example.com</envelope_from>
 <header_from>example.com</header_from>
 </identifiers>
 <auth_results>
 <dkim>
 <domain>example.com</domain>
 <selector>ny</selector>
 <result>fail</result>
 </dkim>
 <dkim>
 <domain>example.net</domain>
 <selector></selector>
 <result>pass</result>
 </dkim>
 <spf>
 <domain>example.com</domain>
 <scope>mfrom</scope>
 <result>pass</result>
 </spf>
 </auth_results>
 </record>
</feedback>

```





## Text Resources

---

- [Overview of Text Resources, page 21-1](#)
- [Content Dictionaries, page 21-2](#)
- [Using and Testing the Content Dictionaries Filter Rules, page 21-6](#)
- [Understanding Text Resources, page 21-8](#)
- [Overview of Text Resource Management, page 21-9](#)
- [Using Text Resources, page 21-12](#)

## Overview of Text Resources

This chapter discusses creating and managing various text resources, such as content dictionaries, disclaimers, and templates.

### Related Topics

- [Content Dictionaries, page 21-1](#)
- [Text Resources, page 21-2](#)
- [Message Disclaimer Stamping, page 21-2](#)
- [Using Custom Dictionaries of Sensitive DLP Terms \(Custom DLP Policies Only\), page 17-17](#)

## Content Dictionaries

You can use content dictionaries to scan messages against message or content filters in order to take appropriate action in accordance with your corporate policies. You can create, delete, and view dictionaries; add and delete entries from a dictionary; and import and export entire dictionaries. You can also determine case sensitivity and word boundary detection for each dictionary. For example, you could create a list of confidential or profane words, and, using a filter rule to scan messages for words in the list, drop or archive messages containing matching words. And you can add a “weight” terms in a dictionary so that certain terms trigger a filter action more easily.

Dictionaries can contain non-ASCII characters.

## Text Resources

Text resources are text objects, such as disclaimers, notification templates, and anti-virus templates. You can create new objects for use in various components of AsyncOS. You can import and export text resources.

## Message Disclaimer Stamping

Message disclaimer stamping allows you to add a disclaimer text resource to messages. For example, you could append a copyright statement, promotional message, or disclaimer to every message sent from within your enterprise.

## Content Dictionaries

Content dictionaries are groups of words or entries that work in conjunction with the Body Scanning feature on the appliance and are available to both content and message filters. Use the dictionaries you define to scan messages, message headers, and message attachments for terms included in the dictionary in order to take appropriate action in accordance with your corporate policies. For example, you could create a list of confidential or profane words, and, using a filter rule to scan messages that contain words in the list, drop, archive, or quarantine the message.

The AsyncOS operating system includes the ability to define a total of 100 content dictionaries using the GUI (Mail Policies > Dictionaries) or the CLI's `dictionaryconfig` command. You can create, delete, and view dictionaries; add and delete entries from a dictionary; and import and export entire dictionaries.

### Related Topics

- [Dictionary Content, page 21-2](#)
- [Importing and Exporting Dictionaries as Text Files, page 21-3](#)
- [Adding Dictionaries, page 21-4](#)
- [Deleting Dictionaries, page 21-5](#)
- [Importing Dictionaries, page 21-5](#)
- [Exporting Dictionaries, page 21-5](#)

## Dictionary Content

Words in dictionaries are created with one text string per line, and entries can be in plain text or in the form of regular expressions. Dictionaries can also contain non-ASCII characters. Defining dictionaries of regular expressions can provide more flexibility in matching terms, but doing so requires you to understand how to delimit words properly. For a more detailed discussion of Python style regular expressions, consult the Python Regular Expression HOWTO, accessible from

<http://www.python.org/doc/howto/>



### Note

To use the special character # at the beginning of a dictionary entry, you can use a character class [#] to prevent it being treated as a comment.

For each term, you specify a “weight,” so that certain terms can trigger filter conditions more easily. When AsyncOS scans messages for the content dictionary terms, it “scores” the message by multiplying the number of term instances by the weight of term. Two instances of a term with a weight of three would result in a score of six. AsyncOS then compares this score with a threshold value associated with the content or message filter to determine if the message should trigger the filter action.

You can also add smart identifiers to a content dictionary. Smart identifiers are algorithms that search for patterns in data that correspond to common numeric patterns, such as social security numbers and ABA routing numbers. These identifiers can be useful for policy enforcement. For more information about regular expressions, see “Regular Expressions in Rules” in the “Using Message Filters to Enforce Email Policies” chapter. For more information about smart identifiers, see “Smart Identifiers” in the “Using Message Filters to Enforce Email Policies” chapter.

**Note**

Dictionaries containing non-ASCII characters may or may not display properly in the CLI on your terminal. The best way to view and change dictionaries that contain non-ASCII characters is to export the dictionary to a text file, edit that text file, and then import the new file back into the appliance. For more information, see [Importing and Exporting Dictionaries as Text Files, page 21-3](#).

**Related Topics**

- [Word Boundaries and Double-byte Character Sets, page 21-3](#)

## Word Boundaries and Double-byte Character Sets

In some languages (double-byte character sets), the concepts of a word or word boundary, or case do not exist. Complex regular expressions that depend on concepts like what is or is not a character that would compose a word (represented as “w” in regex syntax) cause problems when the locale is unknown or if the encoding is not known for certain. For that reason, you may want to disable word-boundary enforcement.

## Importing and Exporting Dictionaries as Text Files

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

- `config.dtd`
- `profanity.txt`
- `proprietary_content.txt`
- `sexual_content.txt`

These text files are intended to be used in conjunction with the content dictionaries feature to aid you in creating new dictionaries. These content dictionaries are weighted and use smart identifiers to better detect patterns in data and trigger filters when the patterns indicate compliance issues.

**Note**

Importing and exporting dictionaries does not preserve the Match Whole Words and Case Sensitive settings. These settings are only preserved in the configuration file.

See [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) for more information accessing on the configuration directory.



You can also create your own dictionary files and import them onto the appliance. The best way to add non-ASCII characters to dictionaries is to add the terms into the dictionary in a text file off the appliance, move that file onto the appliance, and then import that file as a new dictionary. For more information about importing dictionaries, see [Importing Dictionaries, page 21-5](#). For information about exporting dictionaries, see [Exporting Dictionaries, page 21-5](#).

**Caution**

These text files contain terms that some persons may consider obscene, indecent or offensive. If you import terms from these files into your content dictionaries, the terms will be displayed when you later view the content dictionaries you have configured on the appliance.

## Adding Dictionaries

### Procedure

**Step 1** Navigate to the **Mail Policies > Dictionaries** page.

**Step 2** Click **Add Dictionary**.

**Step 3** Type a name for the dictionary.

**Step 4** (Optional) Configure Advanced Matching.

**Note**

AsyncOS preserves the **Match Whole Words** and **Case Sensitive** settings when you save them in the configuration file. AsyncOS does not preserve these settings when importing and exporting dictionaries.

**Step 5** (Optional) Add a smart-identifier to the dictionary.

Smart identifiers are algorithms that search for patterns in data that correspond to common numeric patterns, such as social security numbers and ABA routing numbers. For more information about smart identifiers, see the “Using Message Filters to Enforce Email Policies” chapter.

**Step 6** Enter new dictionary entries into the list of terms.

If you have multiple new entries to add, and you want them to be equally likely trigger a filter action, put each new term on its own line.

**Note**

Content dictionary entries with the regular expression: “. \*” at the beginning or end will cause the system to lock if a match for the “word” MIME part is found. Cisco Systems recommends you do not use “. \*” at the beginning or end of a content dictionary entry.

**Step 7** Specify a weight for the term(s).

You can “weight” a dictionary term so that it is more likely than other terms to trigger a filter action. For more information about how this weight is used to determine filter actions, see “Threshold Scoring for Content Dictionaries” in the “Using Message Filters to Enforce Email Policies” chapter.

**Step 8** Click **Add**.

**Step 9** Submit and commit your changes.

### Related Topics

- [Dictionary Content, page 21-2](#).

## Deleting Dictionaries

### Before You Begin

Be aware that AsyncOS marks any message filter that references the deleted dictionary as invalid. AsyncOS leaves any content filter that references the deleted dictionary enabled, but will evaluate them to false.

### Procedure

- 
- Step 1** Navigate to the **Mail Policies > Dictionaries** page.
  - Step 2** Click the trash can icon next to the dictionary to delete in the dictionary listing.  
A confirmation message lists any filters that are currently referencing the dictionary.
  - Step 3** Click **Delete** in the confirmation message.
  - Step 4** Commit your changes.
- 

## Importing Dictionaries

### Before You Begin

Verify that the file to import is present in the configuration directory on the appliance.

### Procedure

- 
- Step 1** Navigate to the **Mail Policies > Dictionaries** page.
  - Step 2** Click **Import Dictionary**.
  - Step 3** Select the location to import from.
  - Step 4** Select the file to import.
  - Step 5** Select the default weight to use for dictionary terms.  
AsyncOS will assign a default weight to any terms with unspecified weights. You can edit the weights after importing the file.
  - Step 6** Select an encoding.
  - Step 7** Click **Next**.
  - Step 8** Name and edit the dictionary.
  - Step 9** Submit and commit your changes.
- 

## Exporting Dictionaries

### Procedure

- 
- Step 1** Navigate to the **Mail Policies > Dictionaries** page.

- Step 2** Click **Export Dictionary**.
- Step 3** Select the dictionary to export.
- Step 4** Enter a file name for the exported dictionary.  
This is the name of the file that will be created in the configuration directory on the appliance.
- Step 5** Select the location to export to.
- Step 6** Select an encoding for the text file.
- Step 7** Submit and commit your changes.
- 

## Using and Testing the Content Dictionaries Filter Rules

Dictionaries can be used along with the various `dictionary-match()` message filter rules and with content filters.

### Related Topics

- [Dictionary Match Filter Rule, page 21-6](#)

## Dictionary Match Filter Rule

The message filter rule named `dictionary-match(<dictionary_name>)` (and its counterparts) evaluates to true if the message body contains any of the regular expressions in the content dictionary named *dictionary\_name*. If that dictionary does not exist, the rule evaluates to false.

Note that the `dictionary-match()` rule functions similarly to the `body-contains()` body scanning rule: it only scans the body and attachments of messages, and not the headers.

For scanning headers, you can use the appropriate `*-dictionary-match()`-type rule (there are rules for specific headers, such as `subject-dictionary-match()` and a more generic rule, `header-dictionary-match()`, in which you can specify any header including custom headers). See “Dictionary Rules” in the “Using Message Filters to Enforce Email Policies” chapter for more information about dictionary matching.

**Table 21-1**      *Message Filter Rules for Content Dictionaries*

| Rule             | Syntax                                                 | Description                                                                                              |
|------------------|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| Dictionary Match | <code>dictionary-match(&lt;dictionary_name&gt;)</code> | Does the message contain a word that matches all the regular expressions listed in the named dictionary? |

In the following example, a new message filter using the `dictionary-match()` rule is created to blind carbon copy the administrator when the appliance scans a message that contains any words within the dictionary named “secret\_words” (created in the previous example). Note that because of the settings, only messages that contain the whole word “codename” matching the case exactly will evaluate to true for this filter.

```
bcc_codenames:

 if (dictionary-match ('secret_words'))

 {

 bcc('administrator@example.com');

 }
```

In this example, we send the message to the Policy quarantine:

```
quarantine_codenames:

 if (dictionary-match ('secret_words'))

 {

 quarantine('Policy');

 }
```

#### Related Topics

- [Example Dictionary Entries, page 21-7](#)
- [Testing Content Dictionaries, page 21-8](#)

## Example Dictionary Entries

**Table 21-2**      *Example Dictionary Entries*

| Description                                 | Example                                                                                                                                         |
|---------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| Wildcard                                    | *                                                                                                                                               |
| Anchors                                     | Ends with: foo\$<br>Begins with: ^foo                                                                                                           |
| Email address<br>(Do not escape the period) | foo@example.com, @example.com<br>example.com\$ (ends with)<br>@example.*                                                                        |
| Subject                                     | An email subject<br>(keep in mind when using the ^ anchor in email subjects that subjects are often prepended with “RE:” or “FW:” and the like) |

## Testing Content Dictionaries

The `trace` function can provide quick feedback on message filters that use the `dictionary-match()` rule. See [Debugging Mail Flow Using Test Messages: Trace, page 40-1](#) for more information. You can also use the `quarantine()` action to test filters, as in the `quarantine_codenames` filter example above.

## Understanding Text Resources

Text resources are text templates that can be attached to messages or sent as messages. Text resources can be one of the following types:

- **Message disclaimers** — Text that is added to messages. For more information, see [Disclaimer Template, page 21-12](#).
- **Notification templates** — Messages that are sent as notifications, used with the `notify()` and `notify-bcc()` actions. For more information, see [Notification Templates, page 21-19](#).
- **Anti-virus Notification templates** — Messages that are sent as notifications when a virus is found in a message. You can create a template for a container (which appends the original message), or as a notice that is sent without the appended message. For more information, see [Anti-Virus Notification Templates, page 21-20](#).
- **Bounce and Encryption Failure Notification templates** — Messages that are sent as notifications when a message is bounced or message encryption fails. For more information, see [Bounce and Encryption Failure Notification Templates, page 21-23](#).
- **Encryption Notification Templates** — Messages that are sent when you configure the appliance to encrypt outgoing email. The message notifies recipients that they have received an encrypted message and provides instructions for reading it. For more information, see [Encryption Notification Templates, page 21-24](#).

You can use the CLI (`textconfig`) or the GUI to manage text resources, including: adding, deleting, editing, importing, and exporting. For information on managing text resources using the GUI, see [Overview of Text Resource Management, page 21-9](#).

Text resources can contain non-ASCII characters.



### Note

Text resources containing non-ASCII characters may or may not display properly in the CLI on your terminal. To view and change text resources that contain non-ASCII characters, export the text resource to a text file, edit that text file, and then import the new file back into the appliance. For more information, see [Importing and Exporting Text Resources as Text Files, page 21-8](#).

### Related Topics

- [Importing and Exporting Text Resources as Text Files, page 21-8](#)

## Importing and Exporting Text Resources as Text Files

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

See [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) for more information on accessing the configuration directory.

To add non-ASCII characters to text resources, add the terms into the text resource in a text file off the appliance, move that file onto the appliance, and then import that file as a new text resource. For more information about importing text resources, see [Importing Text Resources, page 21-10](#). For information about exporting text resources, see [Exporting Text Resources, page 21-10](#).

## Overview of Text Resource Management

You can manage text resources using either the GUI or the CLI. This section focuses on the GUI.

Manage text resources from the CLI using the `textconfig` command.

Text resource management includes these tasks:

- Adding
- Editing and deleting
- Exporting, and importing
- Defining plain text messages for all text resource types
- Defining HTML-based messages for some text resource types

### Related topics

- [Adding Text Resources, page 21-9](#)
- [Deleting Text Resources, page 21-10](#)
- [Importing Text Resources, page 21-10](#)
- [Exporting Text Resources, page 21-10](#)
- [Overview of HTML-Based Text Resources, page 21-11](#).

## Adding Text Resources

### Procedure

- 
- |               |                                                                                                                                                                                                                                                                                             |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Navigate to <b>Mail Policies &gt; Text Resources</b>                                                                                                                                                                                                                                        |
| <b>Step 2</b> | Click <b>Add Text Resource</b> .                                                                                                                                                                                                                                                            |
| <b>Step 3</b> | Enter a name for the text resource in the <b>Name</b> field.                                                                                                                                                                                                                                |
| <b>Step 4</b> | Select the type of text resource from the <b>Type</b> field.                                                                                                                                                                                                                                |
| <b>Step 5</b> | Enter the message text in either the <b>Text</b> or the <b>HTML and Plain Text</b> field.<br><br>If the text resource allows only plain text messages, use the <b>Text</b> field. If the text resource allows both HTML and plain text messages, use the <b>HTML and Plain Text</b> fields. |
| <b>Step 6</b> | Submit and commit your changes.                                                                                                                                                                                                                                                             |
- 

### Related topics

- [Overview of HTML-Based Text Resources, page 21-11](#).

## Deleting Text Resources

### Before you begin

Note the impact of deleting text resources:

- Any message filters that reference the deleted text resource are marked as invalid.
- Any content filters that reference the deleted text resource are left enabled, but will evaluate to false.

### Procedure

- 
- |               |                                                                                                                                                                                       |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | On the <b>Mail Policies &gt; Text Resources</b> page, click the trash can icon under the Delete column for the text resource you want to delete. A confirmation message is displayed. |
| <b>Step 2</b> | Click <b>Delete</b> to delete the text resource.                                                                                                                                      |
| <b>Step 3</b> | Commit your changes.                                                                                                                                                                  |
- 

## Importing Text Resources

### Before you begin

Ensure that the file to import is in the configuration directory on the appliance.

### Procedure

- 
- |               |                                                                                           |
|---------------|-------------------------------------------------------------------------------------------|
| <b>Step 1</b> | On the <b>Mail Policies &gt; Text Resources</b> page, click <b>Import Text Resource</b> . |
| <b>Step 2</b> | Select a file to import.                                                                  |
| <b>Step 3</b> | Specify an encoding.                                                                      |
| <b>Step 4</b> | Click <b>Next</b> .                                                                       |
| <b>Step 5</b> | Choose a name, edit, and select the text resource type.                                   |
| <b>Step 6</b> | Submit and commit your changes.                                                           |
- 

## Exporting Text Resources

### Before you begin

Be aware that when you export a text resource, a text file is created in the configuration directory on the appliance.

### Procedure

- 
- |               |                                                                                           |
|---------------|-------------------------------------------------------------------------------------------|
| <b>Step 1</b> | On the <b>Mail Policies &gt; Text Resources</b> page, click <b>Export Text Resource</b> . |
| <b>Step 2</b> | Select a text resource to export.                                                         |
| <b>Step 3</b> | Enter a file name for the text resource.                                                  |

- Step 4** Select an encoding for the text file.
- Step 5** Click **Submit** to create the text file containing the text resource in the configuration directory.
- 

## Overview of HTML-Based Text Resources

You can create some text resources with both HTML-based and plain text messages, such as Disclaimers. When a text resource containing both HTML-based and plain text messages is applied to an email message, the HTML-based text resource message is applied to the text/html part of the email message, and the plain text message is applied to the text/plain part of the email message.

When you add or edit an HTML-based text resource, the GUI includes a rich text edit that allows you to enter rich text without having to manually write HTML code.

Consider the following information when adding and editing an HTML-based text resource:

- You can choose to have the plain text version of the message to be automatically generated based on the HTML version, or you can define the plain text version independently.
- You can switch between the rich text editor and HTML code by clicking the **Code View** button.
- To enter HTML code that is not supported in the rich text editor in the GUI, switch to code view and manually enter HTML code. For example, you might want to do this to insert a reference to an image file located on an external server using the `<img src>` HTML tag.

### Related Topics

- [Importing and Exporting HTML-Based Text Resources, page 21-11](#)

## Importing and Exporting HTML-Based Text Resources

You can export to and import from a text file HTML-based text resources. When you export an HTML-based text resource to a file, the file contains the following sections for each version of the text resource:

- [html\_version]
- [text\_version]

The order of these sections does not matter.

For example, an exported file might contain the following text:

```
[html_version]
<p>Sample <i>message.</i></p>
[text_version]
Sample message.
```

Consider the following rules and guidelines when exporting and importing HTML-based text resources:

- When you export an HTML-based text resource whose plain text message is automatically generated from the HTML version, the exported file does not contain the [text\_version] section.
- When you import from a text file, any HTML code under the [html\_version] section is converted to the HTML message in the created text resource if the text resource type supports HTML messages. Similarly, any text under the [text\_version] section is converted to the plain text message in the created text resource.



- When you import from a file that contains an empty or nonexistent `[html_version]` section to create a HTML-based text resource, the appliance creates both an HTML and plain text message using the text in the `[text_version]` section.

## Using Text Resources

All types of text resources are created in the same way, using the Text Resources page or the `textconfig` CLI command. Once created, each type is used in a different way. Disclaimers and notification templates are used with filters and listeners, while anti-virus notification templates are used with mail policies and anti-virus settings.

### Related Topics

- [Disclaimer Template, page 21-12](#)
- [Disclaimer Stamping and Multiple Encodings, page 21-16](#)
- [Notification Templates, page 21-19](#)
- [Anti-Virus Notification Templates, page 21-20](#)
- [Bounce and Encryption Failure Notification Templates, page 21-23](#)
- [Encryption Notification Templates, page 21-24](#)

## Disclaimer Template

The appliance can add a default disclaimer above or below the text (heading or footer) for some or all messages received by a listener. You can add disclaimers to messages on the appliance using the following methods:

- Via a listener, using the GUI or the `listenerconfig` command (see [Adding Disclaimer Text via a Listener, page 21-13](#)).
- Using the content filter action, `Add Disclaimer Text` (see [Content Filter Actions, page 11-9](#)).
- Using the message filter action, `add-footer()` (see the “Using Message Filters to Enforce Email Policies” chapter).
- Using a data loss prevention profile (see [Data Loss Prevention, page 17-1](#)).
- Using message modification for Outbreak Filters to alert the user that the message may be an attempt at phishing or malware distribution (see [Modifying Messages, page 14-6](#)). Disclaimers added for this type of notification are added above the text.

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

Prior to using disclaimer text you have to create the disclaimer template. Use the Text Resources page in the GUI (see [Adding Text Resources, page 21-9](#)) or the `textconfig` command (see the *Cisco AsyncOS CLI Reference Guide*) to create and manage a set of text strings to be used.

### Related Topics

- [Adding Disclaimer Text via a Listener, page 21-13](#)
- [Adding Disclaimers via Filters, page 21-13](#)
- [Disclaimers and Filter Action Variables, page 21-13](#)

## 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 appliance will stamp the disclaimer on both parts of the message. However, if your message has signed content, the content will not be modified because the modification will invalidate the signature. Instead, a new part is created with a disclaimer stamp that says “Content-Disposition inline attachment.” For more information on multipart messages, see “Message Bodies vs. Message Attachments” in the “Using Message Filters to Enforce Email Policies” chapter.

The following example shows how to select a disclaimer to apply to messages on a listener via the GUI:

**Figure 21-1** *Editing a Listener to Include a Disclaimer*  
Add Listener

| Listener Settings               |                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------|
| Name:                           | <input type="text"/>                                                                               |
| Type of Listener:               | <input checked="" type="radio"/> Public<br><input type="radio"/> Private                           |
| Interface:                      | Management TCP Port: 25                                                                            |
| Bounce Profile:                 | Default                                                                                            |
| Disclaimer Above:               | None<br><small>Disclaimer text will be applied above the message body.</small>                     |
| Disclaimer Below:               | None<br><small>Disclaimer text will be applied below the message body.</small>                     |
| SMTP Authentication Profile:    | None                                                                                               |
| Certificate:                    | System Default                                                                                     |
| ▶ SMTP Address Parsing Options: | Optional settings for controlling parsing in SMTP "MAIL FROM" and "RCPT TO"                        |
| ▶ Advanced:                     | Optional settings for customizing the behavior of the Listener                                     |
| ▶ LDAP Queries:                 | No LDAP Server Profiles have been created. Profiles can be defined at System Administration > LDAP |
| SMTP Call-Ahead Profile:        | None                                                                                               |

Cancel Submit

## Adding Disclaimers via Filters

You can also append specific, predefined text strings to the disclaimers of messages using the filter action `add-footer()` or the content filter action “Add Disclaimer Text.” For example, the following message filter rule appends the text string named `legal.disclaimer` to all messages sent from users in the LDAP group “Legal:”

```
Add-Disclaimer-For-Legal-Team:

if (mail-from-group == 'Legal')
{
 add-footer('legal.disclaimer');
}
```

## Disclaimers and Filter Action Variables

You can also use message filter action variables (see “Action Variables” in the “Using Message Filters to Enforce Email Policies” chapter for more information).

The following variables are available for the Disclaimer Template:

**Table 21-3 Anti-Virus Notification Variables**

| Variable                    | Substituted With                                                                                                                                                     |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>\$To</b>                 | Replaced by the message To: header (not the Envelope Recipient).                                                                                                     |
| <b>\$From</b>               | Replaced by the message From: header (not the Envelope Sender).                                                                                                      |
| <b>\$Subject</b>            | Replaced by the subject of the original message.                                                                                                                     |
| <b>\$Date</b>               | Replaced by the current date, using the format MM/DD/YYYY.                                                                                                           |
| <b>\$Time</b>               | Replaced by the current time, in the local time zone.                                                                                                                |
| <b>\$GMTimestamp</b>        | Replaced by the current time and date, as would be found in the Received: line of an email message, using GMT.                                                       |
| <b>\$MID</b>                | 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). |
| <b>\$Group</b>              | 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.      |
| <b>\$Policy</b>             | 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.  |
| <b>\$Reputation</b>         | Replaced by the SenderBase Reputation score of the sender. If there is no reputation score, it is replaced with “None”.                                              |
| <b>\$filenames</b>          | Replaced with a comma-separated list of the message’s attachments’ filenames.                                                                                        |
| <b>\$filetypes</b>          | Replaced with a comma-separated list of the message’s attachments’ file types.                                                                                       |
| <b>\$filesizes</b>          | Replaced with a comma-separated list of the message’s attachment’s file sizes.                                                                                       |
| <b>\$remotehost</b>         | Replaced by the hostname of the system that sent the message to the Email Security appliance.                                                                        |
| <b>\$AllHeaders</b>         | Replaced by the message headers.                                                                                                                                     |
| <b>\$EnvelopeFrom</b>       | Replaced by the Envelope Sender (Envelope From, <MAIL FROM>) of the message.                                                                                         |
| <b>\$Hostname</b>           | Replaced by the hostname of the Email Security appliance.                                                                                                            |
| <b>\$header['string']</b>   | Replaced by the value of the quoted header, if the original message contains a matching header. Note that double quotes may also be used.                            |
| <b>\$enveloperecipients</b> | Replaced by all Envelope Recipients (Envelope To, <RCPT TO>) of the message.                                                                                         |
| <b>\$bodysize</b>           | Replaced by the size, in bytes, of the message.                                                                                                                      |
| <b>\$FilterName</b>         | Returns the name of the filter being processed.                                                                                                                      |

**Table 21-3 Anti-Virus Notification Variables (continued)**

| Variable                    | Substituted With                                                                                                                                                                                                                                                |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>\$MatchedContent</b>     | Returns the content that triggered a scanning filter rule (including filter rules such as <code>body-contains</code> and content dictionaries).                                                                                                                 |
| <b>\$DLPPolicy</b>          | Replaced by the name of the email DLP policy violated.                                                                                                                                                                                                          |
| <b>\$DLPSeverity</b>        | Replaced by the severity of violation. Can be “Low,” “Medium,” “High,” or “Critical.”                                                                                                                                                                           |
| <b>\$DLPRiskFactor</b>      | Replaced by the risk factor of the message’s sensitive material (score 0 - 100).                                                                                                                                                                                |
| <b>\$threat_category</b>    | Replaced with the type of Outbreak Filters threat, such as phishing, virus, scam, or malware.                                                                                                                                                                   |
| <b>\$threat_type</b>        | Replaced by a subcategory of the Outbreak Filters threat category. For example, can be a charity scam, a financial phishing attempt, a fake deal, etc.                                                                                                          |
| <b>\$threat_description</b> | Replaced by a description of the Outbreak Filters threat.                                                                                                                                                                                                       |
| <b>\$threat_level</b>       | Replaced by the message’s threat level (score 0 - 5).                                                                                                                                                                                                           |
| <b>\$threat_verdict</b>     | Replaced by Yes or No, depending on the Message Modification Threat Level threshold. If the viral or non-viral threat level of a message is greater than or equal to the message modification threat level threshold, the value of this variable is set to Yes. |

To use message filter action variables in disclaimers, create a message disclaimer (via the Text Resource page in the GUI or the **textconfig** command), and reference the variable:

```
(running textconfig command)
```

```
Enter or paste the message disclaimer here. Enter '.' on a blank line to end.
```

```
This message processed at: $Timestamp
```

```
.
```

```
Message disclaimer "legal.disclaimervar" created.
```

```
Current Text Resources:
```

1. legal.disclaimer (Message Disclaimer)
2. legal.disclaimervar (Message Disclaimer)

```
Choose the operation you want to perform:
```

```
- NEW - Create a new text resource.

- IMPORT - Import a text resource from a file.

- EXPORT - Export text resource to a file.

- PRINT - Display the content of a resource.

- EDIT - Modify a resource.

- DELETE - Remove a resource from the system.

[]>
```

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

Now, use the new disclaimer in a filter

```
Add-Timestamp:

if (mail-from-group == 'Legal')
{

 add-footer('legal.disclaimervar');

}
```

The `add-footer()` action supports non-ASCII text by adding the footer as an inline, UTF-8 coded, quoted printable attachment.

## Disclaimer Stamping and Multiple Encodings

AsyncOS includes a setting used to modify the way disclaimer stamping with different character encodings works. By default, AsyncOS attempts to place the disclaimers it attaches within the body part of an email message. You can use a setting configured within the `localeconfig` command to configure the behavior if the encodings of the body part and the disclaimer are different. To understand this setting, it is helpful to view an email message as consisting of several parts:

|                                                                               |                       |
|-------------------------------------------------------------------------------|-----------------------|
| To: joe@example.com<br>From: mary@example.com<br>Subject: Hi!<br><blank line> | Headers               |
| Hello!                                                                        | Body part             |
| This message has been scanned...                                              | First attachment part |

|             |                        |
|-------------|------------------------|
|             |                        |
| Example.zip | Second attachment part |

The message body after the first blank line may contain many MIME parts. The second and following parts are often called “attachments,” while the first is often called the “body” or “text.”

A disclaimer can be included in an email as either an attachment (above) or as part of the body

|                                  |                                      |
|----------------------------------|--------------------------------------|
| To: joe@example.com              | Headers                              |
| From: mary@example.com           |                                      |
| Subject: Hi!                     |                                      |
| <blank line>                     |                                      |
| Hello!                           | Body part                            |
| This message has been scanned... | Disclaimer now included in body part |
|                                  |                                      |
| Example.zip                      | First attachment part                |

Typically, when there is an encoding mismatch between the message body and a disclaimer, AsyncOS attempts to encode the entire message in the same encoding as the message body so that the disclaimer will be included in the body (“inline”) and not included as a separate attachment. In other words, the disclaimer will be included inline if the encoding of the disclaimer matches that of the body, or if the text in the disclaimer contains characters that can be displayed inline (in the body). For example, it is possible to have a ISO-8859-1 encoded disclaimer that only contains US-ASCII characters; consequently, this will display “inline” without problems.

However, if the disclaimer cannot be combined with the body, you can use the `localeconfig` command to configure AsyncOS to attempt to promote, or convert, the body text to match the encoding of the disclaimer so that the disclaimer can be included in the body of the message:

```
example.com> localeconfig
```

```
Behavior when modifying headers: Use encoding of message body
```

```
Behavior for untagged non-ASCII headers: Impose encoding of message body
```

```
Behavior for mismatched footer or heading encoding: Only try encoding from
message body
```

```
Choose the operation you want to perform:
```

```
- SETUP - Configure multi-lingual settings.
```

```
[> setup
```

If a header is modified, encode the new header in the same encoding as the message body? (Some MUAs incorrectly handle headers encoded in a different encoding than the body. However, encoding a modified header in the same encoding as the message body may cause certain characters in the modified header to be lost.) [Y]>

If a non-ASCII header is not properly tagged with a character set and is being used or modified, impose the encoding of the body on the header during processing and final representation of the message? (Many MUAs create non-RFC-compliant headers that are then handled in an undefined way. Some MUAs handle headers encoded in character sets that differ from that of the main body in an incorrect way. Imposing the encoding of the body on the header may encode the header more precisely. This will be used to interpret the content of headers for processing, it will not modify or rewrite the header unless that is done explicitly as part of the processing.) [Y]>

Footers or headings are added in-line with the message body whenever possible. However, if the footer or heading is encoded differently than the message body, and if imposing a single encoding will cause loss of characters, it will be added as an attachment. The system will always try to use the message body's encoding for the footer or heading. If that fails, and if the message body's encoding is US-ASCII, the system can try to edit the message body to use the footer's or heading's encoding. Should the system try to impose the footer's or headings's encoding on the message body? [N]> **y**

Behavior when modifying headers: Use encoding of message body  
Behavior for untagged non-ASCII headers: Impose encoding of message body.  
Behavior for mismatched footer or heading encoding: Try both

body and footer or heading encodings

Choose the operation you want to perform:

- SETUP - Configure multi-lingual settings.

For more information about the `localeconfig` command, see the “Configuring the Appliance to Receive Mail” chapter.

## Notification Templates

Notification templates are used with the `notify()` and `notify-copy()` filter actions. Notification templates may contain non-ascii text and action variables (see “Action Variables” in the “Using Message Filters to Enforce Email Policies” chapter), including the anti-virus-related variables used by anti-virus notifications. For example, you could use the `$Allheaders` action variable to include the headers from the original message. You can configure the From: address for notifications, see [Configuring the Return Address for Appliance Generated Messages, page 33-33](#).

Once you have created a notification template, you can refer to it in content and message filters. [Figure 21-2](#) shows a content filter where the `notify-copy()` filter action is set to send the “grape\_text” notification to “grapewatchers@example.com:”



**Figure 21-2** *Notify Example in a Content Filter*  
**Edit Content Filter**

| Edit Filter                                                               |                                                                                                                        |
|---------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------|
| Name:                                                                     | grapecheck                                                                                                             |
| Currently used by policies:                                               | DEFAULT                                                                                                                |
| Description:                                                              | Looking for grapes.                                                                                                    |
| Order:                                                                    | 1                                                                                                                      |
| Apply filter:                                                             | <input checked="" type="radio"/> If one or more conditions match<br><input type="radio"/> Only if ALL conditions match |
| Conditions                                                                |                                                                                                                        |
| Select New Condition... <span>Add Condition</span>                        |                                                                                                                        |
| Condition                                                                 | Delete                                                                                                                 |
| body-contains("grape")                                                    |                                                                                                                        |
| Actions                                                                   |                                                                                                                        |
| Select New Action... <span>Add Action</span>                              |                                                                                                                        |
| Action                                                                    | Delete                                                                                                                 |
| notify-copy ("grapewatchers@example.com", "Found one!", "", "grape_text") |                                                                                                                        |
| <div> <span>Cancel</span> <span>Submit</span> </div>                      |                                                                                                                        |

## Anti-Virus Notification Templates

There are two types of anti-virus notification templates:

- **anti-virus notification template.** The anti-virus notification template is used when the original message is not attached to the virus notification.
- **anti-virus container template.** The container template is used when the original message is sent as an attachment.

Anti-virus notification templates are used in basically the same way as notification templates except that they are used with the anti-virus engine instead of filters. You can specify a custom notification to send while editing a mail policy. You can configure the From: address for anti-virus notifications. For information, see [Configuring the Return Address for Appliance Generated Messages, page 33-33](#).

### Related Topics

- [Custom Anti-Virus Notification Templates, page 21-20](#)

## Custom Anti-Virus Notification Templates

[Figure 21-3](#) shows a mail policy where a custom anti-virus notification is specified.

**Figure 21-3 Anti-Virus Container Template Notification Example in a Mail Policy**

**Virus Infected Messages:**

Action Applied to Message:

Archive Original Message: ☐ No ☒ Yes

Modify Message Subject: ☐ No ☒ Prepend ☐ Append

**[WARNING : VIRUS DETECTED]**

**Advanced**

Add Custom Header to Message: ☒ No ☐ Yes

Header:

Value:

Container Notification:

[Preview Message Body](#)  
(see Mail Policies > Text Resources > Anti-Virus Container Template)

**Related Topics**

- [Anti-Virus Notification Variables, page 21-21](#)

**Anti-Virus Notification Variables**

When creating an anti-virus notification, you can use any of the notification variables listed in [Table 21-4](#):

**Table 21-4 Anti-Virus Notification Variables**

| Variable                     | Substituted With                                                                                                                                                                                                                                   |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>\$To</b>                  | Replaced by the message To: header (not the Envelope Recipient).                                                                                                                                                                                   |
| <b>\$From</b>                | Replaced by the message From: header (not the Envelope Sender).                                                                                                                                                                                    |
| <b>\$Subject</b>             | Replaced by the subject of the original message.                                                                                                                                                                                                   |
| <b>\$AV_VIRUSES</b>          | Replaced by the list of all the viruses found anywhere in the message:<br>“Unix/Apache.Trojan”, “W32/Bagel-F”                                                                                                                                      |
| <b>\$AV_VIRUS_TABLE</b>      | Replaced by the table of MIME-Part/Attachment names and viruses in each part:<br>“HELLO.SCR” : “W32/Bagel-F”<br><unnamed part of the message> : “Unix/Apache.Trojan”                                                                               |
| <b>\$AV_VERDICT</b>          | Replaced by the anti-virus verdict.                                                                                                                                                                                                                |
| <b>\$AV_DROPPED_TABLE</b>    | Replaced by the table of attachments that were dropped. Each row is composed of a part or filename followed by the list of viruses associated with that part:<br>“HELLO.SCR” : “W32/Bagel-F”, “W32/Bagel-d” “Love.SCR” : “Netsky-c”, “W32/Bagel-d” |
| <b>\$AV_REPAIRED_VIRUSES</b> | Replaced by the list of all the viruses found and repaired.                                                                                                                                                                                        |
| <b>\$AV_REPAIRED_TABLE</b>   | Replaced by the table of all parts and viruses found and repaired:<br>“HELLO.SCR” : “W32/Bagel-F”                                                                                                                                                  |
| <b>\$AV_DROPPED_PARTS</b>    | Replaced by the list of filenames that were dropped:<br>“HELLO.SCR”, “CheckThisOut.exe”                                                                                                                                                            |
| <b>\$AV_REPAIRED_PARTS</b>   | Replaced by the list of filenames or parts that were repaired.                                                                                                                                                                                     |

**Table 21-4 Anti-Virus Notification Variables (continued)**

| Variable                      | Substituted With                                                                                                                                                     |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>\$AV_ENCRYPTED_PARTS</b>   | Replaced by the list of filenames or parts that were encrypted.                                                                                                      |
| <b>\$AV_INFECTED_PARTS</b>    | Replaced by a comma-separated list of filenames for the files that contained a virus.                                                                                |
| <b>\$AV_UNSCANNABLE_PARTS</b> | Replaced by the list of filenames or parts that were unscannable.                                                                                                    |
| <b>\$Date</b>                 | Replaced by the current date, using the format MM/DD/YYYY.                                                                                                           |
| <b>\$Time</b>                 | Replaced by the current time, in the local time zone.                                                                                                                |
| <b>\$GMTimestamp</b>          | Replaced by the current time and date, as would be found in the Received: line of an email message, using GMT.                                                       |
| <b>\$MID</b>                  | 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). |
| <b>\$Group</b>                | 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.      |
| <b>\$Policy</b>               | 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.  |
| <b>\$Reputation</b>           | Replaced by the SenderBase Reputation score of the sender. If there is no reputation score, it is replaced with “None”.                                              |
| <b>\$filenames</b>            | Replaced with a comma-separated list of the message’s attachments’ filenames.                                                                                        |
| <b>\$filetypes</b>            | Replaced with a comma-separated list of the message’s attachments’ file types.                                                                                       |
| <b>\$filesizes</b>            | Replaced with a comma-separated list of the message’s attachment’s file sizes.                                                                                       |
| <b>\$remotehost</b>           | Replaced by the hostname of the system that sent the message to the Email Security appliance.                                                                        |
| <b>\$AllHeaders</b>           | Replaced by the message headers.                                                                                                                                     |
| <b>\$EnvelopeFrom</b>         | Replaced by the Envelope Sender (Envelope From, <MAIL FROM>) of the message.                                                                                         |
| <b>\$Hostname</b>             | Replaced by the hostname of the Email Security appliance.                                                                                                            |

**Note**

Variable names are not case-sensitive. For example, specifying “\$to” is equivalent to specifying “\$To” in the text resource. If an “AV\_” variable is empty in the original message, the string <None> is substituted.

After the text resource has been defined, use the **Mail Policies > Incoming/Outgoing Mail Policies > Edit Anti-Virus Settings** page or the `policyconfig -> edit -> antivirus` command to specify that the original message is to be included as an RFC 822 attachment for Repaired, Unscannable, Encrypted, or Virus Positive messages. See [Send custom alert notification \(to recipient only\)](#), page 12-12 for more information.

## Bounce and Encryption Failure Notification Templates

Bounce and encryption failure notification templates are used in basically the same way as notification templates except that they are used with bounce notifications and message encryption failure notifications. You can specify a custom bounce notification to send while editing a bounce profile and a custom message encryption failure notification while editing an encryption profile.

Figure 21-4 shows a bounce notification template specified in a bounce profile.

**Figure 21-4** Bounce Notification Example in a Bounce Profile



**Note**

You must use RFC-1891 DSNs to use custom templates.

Figure 21-5 shows an encryption failure template specified in an encryption profile.

**Figure 21-5** Encryption Failure Notification Example in an Encryption Profile

### Related Topics

- [Bounce and Encryption Failure Notification Variables, page 21-23](#)

## Bounce and Encryption Failure Notification Variables

When creating a bounce or encryption failure notification, you can use any of the notification variables listed in Table 21-5:

**Table 21-5** Bounce Notification Variables

| Variable              | Substituted With                                                                                               |
|-----------------------|----------------------------------------------------------------------------------------------------------------|
| <b>\$Subject</b>      | The subject of the original message.                                                                           |
| <b>\$Date</b>         | Replaced by the current date, using the format MM/DD/YYYY.                                                     |
| <b>\$Time</b>         | Replaced by the current time, in the local time zone.                                                          |
| <b>\$GMTTimeStamp</b> | Replaced by the current time and date, as would be found in the Received: line of an email message, using GMT. |

**Table 21-5**      **Bounce Notification Variables (continued)**

| Variable                  | Substituted With                                                                                                                                                     |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>\$MID</b>              | 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). |
| <b>\$BouncedRecipient</b> | Bounced recipient address                                                                                                                                            |
| <b>\$BounceReason</b>     | Reason for this notification                                                                                                                                         |
| <b>\$remotehost</b>       | Replaced by the hostname of the system that sent the message to the Email Security appliance.                                                                        |

## Encryption Notification Templates

Encryption notification templates are used when you configure Cisco Email Encryption to encrypt outbound email. The notification informs recipients that they have received an encrypted message and provides instructions for reading it. You can specify a custom encryption notification to send with encrypted messages. You specify both an HTML and a text encryption notification when you create an encryption profile. Therefore, if you want to create a custom profile, you should create both text and HTML notifications.

Figure 21-6 shows encryption notifications specified in an encryption profile.

**Figure 21-6**      **Encryption Notification Templates Enabled on Encryption Profile**

**Notification Settings**

Select a notification template for each format. The notification informs recipients that they have received an encrypted message and provides instructions for reading it.

HTML Notification: encrypt\_html Preview Message

Text Notification: encrypt\_txt Preview Message

Notification templates can be configured in Mail Policies > Text Resources



## Validating Recipients Using an SMTP Server

---

- [Overview of SMTP Call-Ahead Recipient Validation, page 22-1](#)
- [SMTP Call-Ahead Recipient Validation Workflow, page 22-1](#)
- [How to Validate Recipients Using an External SMTP Server, page 22-3](#)
- [Enabling a Listener to Validate Incoming Mail Via the SMTP Server, page 22-6](#)
- [Configuring LDAP Routing Query Settings, page 22-6](#)
- [SMTP Call-Ahead Query Routing, page 22-7](#)
- [Bypassing SMTP Call-Ahead Validation for Certain Users or Groups, page 22-8](#)

### Overview of SMTP Call-Ahead Recipient Validation

The SMTP call-ahead recipient validation feature queries an external SMTP server before accepting incoming mail for a recipient. Use this feature to validate recipients when you cannot use LDAP Accept or the Recipient Access Table (RAT). For example, suppose you host mail for many mailboxes, each using a separate domain, and your LDAP infrastructure does not allow you to query the LDAP server to validate each recipient. In this case, the Email Security appliance can query the SMTP server and validate the recipient before continuing the SMTP conversation.


You can use SMTP call-ahead recipient validation in order to reduce processing on messages for invalid recipients. Typically, a message for an invalid recipient progresses through the work queue before it can be dropped. Instead, an invalid message can be dropped or bounced during the incoming/receiving part of the email pipeline without requiring additional processing.

### SMTP Call-Ahead Recipient Validation Workflow

When you configure your Email Security appliance for SMTP call-ahead recipient validation, the Email Security appliance suspends the SMTP conversation with the sending MTA while it “calls ahead” to the SMTP server to verify the recipient. When the appliance queries the SMTP server, it returns the SMTP server’s response to the Email Security appliance, and depending on the settings you have configured, you can accept the mail or drop the connection with a code and custom response.

[Figure 22-1](#) shows the basic workflow of the SMTP call-head validation conversation.

Figure 22-1 SMTP Call Ahead Server Conversation Workflow



1. The sending MTA initiates an SMTP conversation.
2. The Email Security appliance suspends the SMTP conversation while it sends a query to the SMTP server to verify the recipient, *validuser@recipient.com*.



**Note** If SMTP routes or LDAP routing queries are configured, these routes will be used to query the SMTP server.

3. The SMTP Server returns a query response to the Email Security appliance.
4. The Email Security appliance resumes the SMTP conversation and sends a response to the sending MTA, allowing the conversation to continue or dropping the connection based on the SMTP server response (and settings you configure in the SMTP Call-Ahead profile).

Due to the order of processes in the email pipeline, if the message for a given recipient is rejected by the RAT, then the SMTP call-ahead recipient validation will not occur. For example, if you specified in the RAT that only mail for *example.com* is accepted, then mail for *recipient@domain2.com* is rejected before SMTP call-ahead recipient validation can occur.



**Note**

If you have configured Directory Harvest Attack Prevention (DHAP) in the HAT, be aware that SMTP call-ahead server rejections are part of the number of rejections included in the maximum invalid recipients per hour that you specify. You may need to adjust this number to account for additional SMTP server rejections. For more information about DHAP, see the “Configuring the Gateway to Receive Email” chapter.

# How to Validate Recipients Using an External SMTP Server

**Table 22-1**      *How to Validate Recipients Using an External SMTP Server*

|               | Do This                                                                                                              | More Info                                                                                              |
|---------------|----------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Determine how the appliance connects to the SMTP server and interprets the server's responses.                       | <a href="#">Configuring the Call-Ahead Server Profile, page 22-3</a>                                   |
| <b>Step 2</b> | Configure a public listener to use the SMTP server to validate recipients                                            | <a href="#">Enabling a Listener to Validate Incoming Mail Using an External SMTP Server, page 22-6</a> |
| <b>Step 3</b> | (Optional) Update your LDAP Routing query to determine the SMTP server to use when routing mail to a different host. | <a href="#">Configuring LDAP Routing Query Settings, page 22-7</a>                                     |
| <b>Step 4</b> | (Optional) Configure the appliance to bypass call-ahead validation for certain recipients                            | <a href="#">Bypassing SMTP Call-Ahead Validation for Certain Recipients or Groups, page 22-8</a>       |

## Related Topics

- [Configuring the Call-Ahead Server Profile, page 22-3](#)

## Configuring the Call-Ahead Server Profile

When you configure the SMTP Call-Ahead Server Profile, you specify the settings that determine how the Email Security appliance connects with the SMTP server and how it interprets the responses sent back from the SMTP server.

### Procedure

- 
- |               |                                                                                                                                                                    |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Click <b>Network &gt; SMTP Call-Ahead</b> .                                                                                                                        |
| <b>Step 2</b> | Click <b>Add Profile</b> .                                                                                                                                         |
| <b>Step 3</b> | Enter the settings for the profile. For more information, see <a href="#">Table 22-2SMTP Call-Ahead Server Profile Settings, page 22-4</a> .                       |
| <b>Step 4</b> | Configure the advanced settings for the profile. For more information, see <a href="#">Table 22-3SMTP Call-Ahead Server Profile Advanced Settings, page 22-5</a> . |
| <b>Step 5</b> | Submit and commit your changes.                                                                                                                                    |
- 

## Related Topics

- [SMTP Call-Ahead Server Profile Settings, page 22-3](#)
- [Call Ahead Server Responses, page 22-5](#)

## SMTP Call-Ahead Server Profile Settings

When you configure the SMTP Call-Ahead Server Profile, you need to configure settings that determine how the Email Security appliance connects with the SMTP server.



Table 22-2 SMTP Call-Ahead Server Profile Settings

| Setting                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Profile Name              | Name of the call-ahead server profile.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Call-Ahead Server Type    | <p>Choose from one of the following methods for connecting to the call-ahead server:</p> <ul style="list-style-type: none"> <li>• <b>Use Delivery Host.</b> Select this option to specify that the host for the delivery email address is used for the SMTP call-ahead query. For example, if the mail recipient address is <i>recipient@example.com</i>, the SMTP query is executed against the SMTP server associated with <i>example.com</i>. If you have configured SMTP routes or LDAP routing queries, these routes are used to determine the SMTP server to query. For details about configuring LDAP routing queries, see <a href="#">Configuring LDAP Routing Query Settings, page 22-6</a>.</li> <li>• <b>Static Call-Ahead Server.</b> Use this option to create a static list of call-ahead servers to query. You may want to use this option if you do not expect the names and locations of the call-ahead servers to change often. When you use this option, the Email Security appliance queries the hosts in a round-robin fashion, starting with the first static call-ahead server listed.</li> </ul> <p> <b>Note</b> Note that when you choose the static call-ahead server type, no SMTP routes are applied to the query. Instead an MX lookup is performed, and then an A lookup is performed on the hosts to obtain the call-ahead IP addresses for the static servers.</p> |
| Static Call-Ahead Servers | <p>If you choose to use the static call-ahead server type, enter a list of host and port combinations in this field. List the server and port using the following syntax:</p> <pre>ironport.com:25</pre> <p>Separate multiple entries with a comma.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

Table 22-3 describes the SMTP Call-Ahead Server Profile advanced settings:

**Table 22-3 SMTP Call-Ahead Server Profile Advanced Settings**

| Setting                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface</b>                   | The interface used to initiate the SMTP conversation with the SMTP server.<br>Choose to use the Management interface or Auto. When you select Auto, the Email Security appliance attempts to automatically detect an interface to use. The Cisco IronPort interface attempts to connect to the SMTP server in the following ways: <ul style="list-style-type: none"> <li>• If the call-ahead server is on the same subnet as one of the configured interfaces, then the connection is initiated by the matching interface.</li> <li>• Any configured SMTP routes are used to route the query.</li> <li>• Otherwise, the interface that is on the same subnet as the default gateway are used.</li> </ul> |
| <b>MAIL FROM Address</b>           | The MAIL FROM: address to be used for the SMTP conversation with the SMTP server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Validation Request Timeout</b>  | The number of seconds to wait for a result from the SMTP server. This timeout value is for a single recipient validation request which may involve contacting multiple call-ahead servers. See <a href="#">Call Ahead Server Responses, page 22-5</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Validation Failure Action</b>   | The action to be taken when a recipient validation request fails (due to a timeout, server failure, network issue, or unknown response). You can configure how you want the Email Security appliance to handle the different responses. See <a href="#">Call Ahead Server Responses, page 22-5</a> .                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Temporary Failure Action</b>    | The action to be taken when a recipient validation request temporarily fails (and a 4xx response is returned from the remote SMTP server). This can occur when the mailbox is full, the mailbox is not available, or the service is not available).<br><br>See <a href="#">Call Ahead Server Responses, page 22-5</a> .                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Max. Recipients per Session</b> | Maximum number of recipients to be validated in a single SMTP session.<br>Specify between 1 - 25,000 sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Max. Connections per Server</b> | Maximum number of connections to a single call-ahead SMTP server.<br>Specify between 1-100 connections.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Cache</b>                       | Size of the cache for SMTP responses. Specify between 100-1,000,000 entries                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Cache TTL</b>                   | Time-to-live value for entries in the cache. This field defaults to 900 seconds. Specify between 60 - 86400 seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Call Ahead Server Responses

The SMTP server may return the following responses:

- **2xx:** When an SMTP code starting with 2 is received from the call-ahead server, the recipient is accepted. For example, a response of 250 allows the mailing action to continue.

- **4xx:** An SMTP code starting with a 4 means that a temporary failure has occurred in processing the SMTP request. A retry may later be processed successfully. For example, a response of 451 means the requested action was aborted or there was a local error in processing.
- **5xx:** An SMTP code starting with 5 means a permanent failure in processing the SMTP request occurred. For example, a response of 550 means the requested action was not taken or the mailbox was unavailable.
- **Timeout.** If no response is returned from the call-ahead server, you can configure how long to attempt to retry before a timeout occurs.
- **Connection error.** If a connection to the call-ahead server fails, you can configure whether to accept or reject a connection for the recipient address.

## Enabling a Listener to Validate Incoming Mail Via the SMTP Server

Once you create the SMTP Call-Ahead Server Profile, you need to enable it on a listener to allow the listener to validate incoming mail via the SMTP server. SMTP call-ahead functionality is only available on public listeners, as recipient validation is not necessary for private listeners.

### Procedure

- 
- Step 1** Go to **Network > Listeners**.
  - Step 2** Click the name of the listener where you want to enable SMTP call-ahead functionality.
  - Step 3** In the **SMTP Call Ahead Profile** field, select the SMTP Call-Ahead profile you want to enable.
  - Step 4** Submit and commit your changes.
- 

## Configuring LDAP Routing Query Settings

If you use an LDAP routing query to route mail to a different mail host, AsyncOS uses the Alternate Mailhost Attribute to determine the SMTP server to query. However, there are cases where you may not want that to occur. For example, in the following schema, note that the mail host attribute (`mailHost`) has a different SMTP address than the servers listed in the call-ahead SMTP server attribute (`callAhead`):

```
dn: mail=cisco.com, ou=domains
mail: cisco.com
mailHost: smtp.mydomain.com
policy: ASAV
callAhead: smtp2.mydomain.com,smtp3.mydomain.com:9025
```

In this case, you can use the **SMTP Call-Ahead** field to create a routing query that directs the SMTP call-ahead query to the servers listed in the `callAhead` attribute. For example, you might create a routing query with the following attributes:

**Figure 22-2 LDAP Routing Query Configured for SMTP Call-Ahead:**

|                                                                                                                            |                                       |
|----------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
| <input checked="" type="checkbox"/> Routing Query                                                                          |                                       |
| Name:                                                                                                                      | LDAP1.routing                         |
| Query String:                                                                                                              | {mail={d}} <a href="#">Test Query</a> |
| Recipient Email to Rewrite the Envelope Recipient:                                                                         |                                       |
| Alternative Mailhost Attribute:                                                                                            | mailHost                              |
| SMTP Call-Ahead Server Attribute (optional):                                                                               | callAhead                             |
| <small>This attribute is used only if an SMTP Call-Ahead server is configured. Go to Network &gt; SMTP Call-Ahead.</small> |                                       |

In this query, the {d} represents the domain part of the recipient address, and the SMTP Call-Ahead Server Attribute returns the values for the call-ahead servers and the port that should be used for the query: smtp2.mydomain.com, smtp3.mydomain.com on port 9025.

**Note**

This example shows just one way to configure a query that enables you to use the LDAP routing query to direct SMTP call-ahead queries to the correct SMTP servers. You are not required to use the query string or specific LDAP attributes described in this example.

## SMTP Call-Ahead Query Routing

When routing an SMTP call-ahead query, AsyncOS checks for information in the following order:

**Figure 22-3 SMTP Call Ahead Query Routing Workflow**

Checks the domain name.



Checks for LDAP Routing queries.



Checks for SMTP Routes.



Performs a DNS Lookup (First an MX Lookup is performed, followed by an A lookup).

If there is no LDAP routing query or no SMTP Routes configured for the domain, the result of preceding state is passed to next stage. In any case where there is no SMTP Route present, a DNS lookup is performed.

When you use an LDAP Routing query for an SMTP call-ahead query and you also have SMTP routes configured, the routing behavior depends upon the values returned by the routing query.

- If the LDAP routing query returns a single hostname without a port, the SMTP call-ahead query applies SMTP routes. If the SMTP routes only lists the destination host as the hostname, a DNS lookup is performed to obtain the IP address of the SMTP server.

- If the LDAP routing query returns a single hostname with a port, the SMTP route is used, but the port returned by the LDAP query is used over any ports specified in SMTP routes. If the SMTP routes only lists the destination host as the hostname, a DNS lookup is performed to obtain the IP address of the SMTP server.
- If the LDAP routing query returns multiple hosts with or without ports, SMTP routes are applied, but the ports returned by the LDAP routing query are used over those present in SMTP routes. If the SMTP routes only lists the destination host as the hostname, a DNS lookup is performed to obtain the IP address of the SMTP server.

## Bypassing SMTP Call-Ahead Validation for Certain Users or Groups

You may want to enable SMTP call-ahead validation on a listener but skip the SMTP call-ahead validation for certain users or groups of users.

You may want to skip SMTP call-ahead validation for recipients for whom mail should not be delayed during SMTP call-ahead queries. For example, you could add a RAT entry for a customer service alias that you know is valid and will likely need immediate attention.

To configure bypassing SMTP call-ahead validation via the GUI, select **Bypass SMTP Call-Ahead** when you add or edit the RAT entry.



## Encrypting Communication with Other MTAs

---

- [Overview of Encrypting Communication with Other MTAs, page 23-1](#)
- [Obtaining Certificates, page 23-2](#)
- [Enabling TLS on a Listener's HAT, page 23-6](#)
- [Enabling TLS and Certificate Verification on Delivery, page 23-10](#)
- [Managing Lists of Certificate Authorities, page 23-16](#)
- [Enabling a Certificate for HTTPS, page 23-18](#)

### Overview of Encrypting Communication with Other MTAs

Enterprise Gateways (or Message Transfer Agents, i.e. MTAs) normally communicate “in the clear” over the Internet. That is, the communications are not encrypted. In several scenarios, malicious agents can intercept this communication without the knowledge of the sender or the receiver. Communications can be monitored and even altered by a third party.

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. AsyncOS supports the STARTTLS extension to SMTP (Secure SMTP over TLS), described in RFC 3207 (which obsoletes RFC 2487).

The TLS implementation in AsyncOS provides privacy through encryption. It allows you to import an X.509 certificate and private key from a certificate authority service or create a self-signed certificate to use on the appliance. AsyncOS supports separate TLS certificates for public and private listeners, secure HTTP (HTTPS) management access on an interface, the LDAP interface, and all outgoing TLS connections.

#### Related Topics

- [How to Encrypt SMTP Conversations using TLS, page 23-2](#)

## How to Encrypt SMTP Conversations using TLS

**Table 23-1**      *How to Encrypt SMTP Conversations using TLS*

|               | Do This                                                                                                                                                                      | More Info                                                                                                                                                                                                                                    |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Obtain an X.509 certificate and private key from a recognized certificate authority.                                                                                         | <a href="#">Obtaining Certificates, page 23-2</a>                                                                                                                                                                                            |
| <b>Step 2</b> | Install the certificate on the Email Security appliance                                                                                                                      | Install a certificate by either: <ul style="list-style-type: none"> <li>• <a href="#">Creating a Self-Signed Certificate using the GUI, page 23-3</a></li> <li>• <a href="#">Importing a Certificate Using the GUI, page 23-5</a></li> </ul> |
| <b>Step 3</b> | Enable TLS for receiving messages, delivering messages, or both                                                                                                              | <ul style="list-style-type: none"> <li>• <a href="#">Enabling TLS on a Listener's HAT, page 23-6</a></li> <li>• <a href="#">Enabling TLS and Certificate Verification on Delivery, page 23-10</a></li> </ul>                                 |
| <b>Step 4</b> | (Optional) Customize the list of trusted certificate authorities that the appliance uses to verify a certificate from a remote domain to establish the domain's credentials. | <a href="#">Managing Lists of Certificate Authorities, page 23-16</a>                                                                                                                                                                        |
| <b>Step 5</b> | (Optional) Configure the Email Security appliance to send an alert when it's unable to deliver messages to a domain that requires a TLS connection.                          | <a href="#">Sending Alerts When a Required TLS Connection Fails, page 23-11</a>                                                                                                                                                              |

## Obtaining Certificates

To use TLS, the Email Security appliance must have an X.509 certificate and matching private key for receiving and delivery. You may use the same certificate for both SMTP receiving and delivery and different certificates for HTTPS services on an interface, the LDAP interface, and all outgoing TLS connections to destination domains, or use one certificate for all of them.

You may purchase certificates and private keys from a recognized certificate authority service. A certificate authority is a third-party organization or company that issues digital certificates used to verify identity and distributes public keys. This provides an additional level of assurance that the certificate is issued by a valid and trusted identity. Cisco does not recommend one service over another.

The Email Security appliance can create a self-signed certificate for your own use and generate a Certificate Signing Request (CSR) to submit to a certificate authority to obtain the public certificate. The certificate authority will return a trusted public certificate signed by a private key. Use the **Network > Certificates** page in the GUI or the `certconfig` command in the CLI to create the self-signed certificate, generate the CSR, and install the trusted public certificate.

If you are acquiring or creating a certificate for the first time, search the Internet for “certificate authority services SSL Server Certificates,” and choose the service that best meets the needs of your organization. Follow the service's instructions for obtaining a certificate.

You can view the entire list of certificates on the **Network > Certificates** page in the GUI and in the CLI by using the `print` command after you configure the certificates using `certconfig`. Note that the `print` command does not display intermediate certificates.

**Caution**

Your appliance ships with a demonstration certificate to test the TLS and HTTPS functionality, but enabling either service with the demonstration certificate is not secure and is not recommended for general use. When you enable either service with the default demonstration certificate, a warning message is printed in the CLI.

**Related Topics**

- [Intermediate Certificates, page 23-3](#)
- [Certificates and Centralized Management, page 23-3](#)
- [Creating a Self-Signed Certificate using the GUI, page 23-3](#)
- [Importing a Certificate Using the GUI, page 23-5](#)
- [Creating a Self-Signed Certificate or Importing a Certificate using the CLI, page 23-6](#)
- [Exporting a Certificate Using the GUI, page 23-6](#)

## Intermediate Certificates

In addition to root certificate verification, AsyncOS supports the use of intermediate certificate verification. Intermediate certificates are certificates issued by a trusted root certificate authority which are then used to create additional certificates - effectively creating a chained line of trust. For example, a certificate may be issued by godaddy.com who, in turn, is granted the rights to issue certificates by a trusted root certificate authority. The certificate issued by godaddy.com must be validated against godaddy.com's private key as well as the trusted root certificate authority's private key.

## Certificates and Centralized Management

A certificate usually uses the local machine's hostname for the certificate's common name. If your Email Security appliances are part of a cluster, you will need to import a certificate for each cluster member as the machine level, with the exception of a wild card certificate that you can install at the cluster level. Each cluster member's certificate must use the same certificate name so the cluster can refer to it when a member's listener is communicating with another machine.

## Creating a Self-Signed Certificate using the GUI

You might want to create or import a certificate on the appliance for any of the following reasons:

- To encrypt SMTP conversations with other MTAs using TLS (both inbound and outbound conversations).
- To enable the HTTPS service on the appliance for accessing the GUI using HTTPS.
- Use as a client certificate for LDAPS if the LDAP server asks for a client certificate.
- To allow secure communication between the appliance and RSA Enterprise Manager for DLP.

**Procedure**

- Step 1** Navigate to the Network > Certificates page.



**Step 2** Click **Add Certificate**.

**Step 3** Select **Create Self-Signed Certificate**.

Figure 23-1 shows the Add Certificate page with the Create Self-Signed Certificate option selected.

**Figure 23-1 Add Certificate Page**

#### Add Certificate

| Add Certificate                                                                |                                                                     |
|--------------------------------------------------------------------------------|---------------------------------------------------------------------|
| Add Certificate:                                                               | Create Self-Signed Certificate ▼                                    |
| Common Name:                                                                   | <input type="text"/>                                                |
| Organization:                                                                  | <input type="text"/>                                                |
| Organizational Unit:                                                           | <input type="text"/>                                                |
| City (Locality):                                                               | <input type="text"/>                                                |
| State (Province):                                                              | <input type="text"/>                                                |
| Country:                                                                       | <input type="text"/>                                                |
| Duration before expiration:                                                    | 3650 days                                                           |
| Private Key Size:                                                              | <input checked="" type="radio"/> 2048<br><input type="radio"/> 1024 |
| <input type="button" value="Cancel"/> <input type="button" value="Next &gt;"/> |                                                                     |

**Step 4** Enter the following information for the self-signed certificate:

|                                   |                                                                                            |
|-----------------------------------|--------------------------------------------------------------------------------------------|
| <b>Common Name</b>                | The fully qualified domain name.                                                           |
| <b>Organization</b>               | The exact legal name of the organization.                                                  |
| <b>Organizational Unit</b>        | Section of the organization.                                                               |
| <b>City (Locality)</b>            | The city where the organization is legally located.                                        |
| <b>State (Province)</b>           | The state, county, or region where the organization is legally located.                    |
| <b>Country</b>                    | The two letter ISO abbreviation of the country where the organization is legally located.  |
| <b>Duration before expiration</b> | The number of days before the certificate expires.                                         |
| <b>Private Key Size</b>           | Size of the private key to generate for the CSR. Only 2048-bit and 1024-bit are supported. |

**Step 5** Click **Next** to view the certificate and signature information.

Figure 23-2 shows an example of a self-signed certificate.

**Figure 23-2 View Certificate Page****View Certificate example.com**

| Add Certificate                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Certificate Name:                     | example.com                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Common Name:                          | example.com                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Organization:                         | Example                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Organization Unit:                    | Org                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| City (Locality):                      | San Francisco                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| State (Province):                     | CA                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Country:                              | US                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Signature Issued By:                  | Common Name (CN): example.com<br>Organization (O): Example<br>Organizational Unit (OU): Org<br>Issued On: Feb 17 21:45:33 2010 GMT<br>Expires On: Feb 15 21:45:33 2020 GMT<br><br><i>If you would like a signed certificate, Download the certificate request. Submit this to a certificate authority. Once you receive the signed certificate, Upload it below.</i><br><br>Upload Signed Certificate: <input type="button" value="Browse..."/> <a href="#">Download Certificate Signing Request...</a><br><br><i>Uploading a new certificate will overwrite the existing certificate.</i><br><br><input type="button" value="Cancel"/> <input type="button" value="Submit"/> |
| Intermediate Certificates (optional): | <input type="button" value="Browse..."/> <a href="#">Download Intermediate Certificate Signing Request...</a><br><br><i>Uploading intermediate certificates if applicable.</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

- Step 6** Enter a name for the certificate. AsyncOS assigns the common name previously entered by default.
- Step 7** If you want to submit a CSR for the self-signed certificate to a certificate authority, click **Download Certificate Signing Request** to save the CSR in PEM format to a local or network machine.
- Step 8** Submit and commit your changes.

When the certificate authority returns the trusted public certificate signed by a private key, upload it by clicking on the certificate's name on the Certificates page and entering the path to the file on your local machine or network. Make sure that the trusted public certificate that you receive is in PEM format or a format that you can convert to PEM using before uploading to the appliance. (Tools for doing this are included with OpenSSL, free software from <http://www.openssl.org>.)

Uploading the certificate from the certificate authority overwrites the existing certificate. You can also upload an intermediate certificate related to the self-signed certificate. You can use the certificate with a public or private listener, an IP interface's HTTPS services, the LDAP interface, or all outgoing TLS connections to destination domains.

## Importing a Certificate Using the GUI

AsyncOS also allows you to import certificates saved in the PKCS #12 format to use on your appliance.

### Procedure

- Step 1** Navigate to the Network > Certificates page.
- Step 2** Click **Add Certificate**.
- Step 3** Select the **Import Certificate** option.
- Step 4** Enter the path to the certificate file on your network or local machine.
- Step 5** Enter the password for the file.
- Step 6** Click **Next** to view the certificate's information.
- Step 7** Enter a name for the certificate.

AsyncOS assigns the common name by default.

**Step 8** Submit and commit your changes.

---

## Creating a Self-Signed Certificate or Importing a Certificate using the CLI

To create a self-signed certificate or import a certificate using the CLI, use the `certconfig` command.

## Exporting a Certificate Using the GUI

AsyncOS also allows you to export certificates and save them in the PKCS #12 format.

### Procedure

---

- Step 1** Navigate to the Network > Certificates page.
  - Step 2** Click **Export Certificate**.
  - Step 3** Select the certificate you want to export.
  - Step 4** Enter the file name for the certificate.
  - Step 5** Enter a password for the certificate file.
  - Step 6** Click **Export**.
  - Step 7** Save the file to a local or network machine.
  - Step 8** You can export additional certificates or click **Cancel** to return to the Network > Certificates page.
- 

## Enabling TLS on a Listener's HAT

You must enable TLS for any listeners where you require encryption. You may want to enable TLS on listeners facing the Internet (that is, public listeners), but not for listeners for internal systems (that is, private listeners). Or, you may want to enable encryption for all listeners.

You can specify the following settings for TLS on a listener.

**Table 23-2** *TLS Settings for a Listener*

| TLS Setting | Meaning                                                                                                                                                                                              |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1. No       | TLS is not allowed for incoming connections. No connections to the listener will require encrypted SMTP conversations. This is the default setting for all listeners you configure on the appliance. |

**Table 23-2**      **TLS Settings for a Listener**

| TLS Setting  | Meaning                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2. Preferred | TLS is allowed for incoming connections to the listener from MTAs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 3. Required  | TLS is allowed for incoming connections to the listener from MTAs, and until a STARTTLS command is received, the appliance responds with an error message to every command other than NOOP, EHLO, or QUIT. This behavior is specified by RFC 3207, which defines the SMTP Service Extension for Secure SMTP over Transport Layer Security. “Requiring” TLS means that email which the sender is not willing to encrypt with TLS will be refused by the appliance before it is sent, thereby preventing it from be transmitted in the clear. |

By default, neither private nor public listeners allow TLS connections. You must enable TLS in a listener's HAT to enable TLS for either inbound (receiving) or outbound (sending) email. In addition, all default mail flow policy settings for private and public listeners have the `tls` setting set to “off.”

You can assign a specific certificate for TLS connections to individual public listeners when creating a listener. For more information, see [Listening for Connection Requests by Creating a Listener via the GUI](#), page 5-8.

#### Related Topics

- [Assigning a Certificate to a Public or Private Listener for TLS Connections Using the GUI](#), page 23-7
- [Assigning a Certificate to a Public or Private Listener for TLS Connections Using the CLI](#), page 23-8
- [Logging](#), page 23-8
- [GUI Example: Changing the TLS Setting for Listener's HAT](#), page 23-8
- [CLI Example: Changing the TLS Setting for Listener's HAT](#), page 23-9

## Assigning a Certificate to a Public or Private Listener for TLS Connections Using the GUI

#### Procedure

- 
- |               |                                                 |
|---------------|-------------------------------------------------|
| <b>Step 1</b> | Navigate to the Network > Listeners page.       |
| <b>Step 2</b> | Click the name of the Listener to edit.         |
| <b>Step 3</b> | In the Certificate field, choose a certificate. |
| <b>Step 4</b> | Submit and commit your changes.                 |
-

## Assigning a Certificate to a Public or Private Listener for TLS Connections Using the CLI

### Procedure

- 
- Step 1** Use the `listenerconfig -> edit` command to choose a listener you want to configure.
- Step 2** Use the `certificate` command to see the available certificates.
- Step 3** Choose the certificate you want to assign to the listener when prompted.
- Step 4** When you are finished configuring the listener, issue the `commit` command to enable the change.
- 

## Logging

The Email Security appliance will note in the mail logs instances when TLS is required but could not be used by the listener. The mail logs will be updated when the following conditions are met:

- TLS is set to “required” for a listener.
- The Email Security appliance has sent a “Must issue a STARTTLS command first” command.
- The connection is closed without having received any successful recipients.

Information on why the TLS connection failed will be included in the mail logs.

## GUI Example: Changing the TLS Setting for Listener's HAT

### Procedure

- 
- Step 1** Navigate to the Mail Policies > Mail Flow Policies page.
- Step 2** Choose a listener whose policies you want to modify, and then click the link for the name of policy to edit. (You can also edit the Default Policy Parameters.)
- Step 3** In the “Encryption and Authentication” section, for the “TLS:” field, choose the level of TLS you want for the listener.

**Figure 23-3** Requiring TLS in a Listener's Mail Flow Policy Parameters

|                                |                                                  |                                                                                                     |
|--------------------------------|--------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| Encryption and Authentication: | TLS:                                             | <input checked="" type="radio"/> Off <input type="radio"/> Preferred <input type="radio"/> Required |
|                                | SMTP Authentication:                             | <input checked="" type="radio"/> Off <input type="radio"/> Preferred <input type="radio"/> Required |
|                                | If Both TLS and SMTP Authentication are enabled: | <input type="checkbox"/> Require TLS To Offer SMTP Authentication                                   |

- Step 4** Submit and commit your changes.

The mail flow policy for the listener is updated with the TLS setting you chose.

---

## CLI Example: Changing the TLS Setting for Listener's HAT

### Procedure

- 
- Step 1** Use the `listenerconfig -> edit` command to choose a listener you want to configure.
- Step 2** Use the `hostaccess -> default` command to edit the listener's default HAT settings.
- Step 3** Change the TLS setting by entering one of the following choices when you are prompted with the following questions:

```
Do you want to allow encrypted TLS connections?
```

```
1. No
```

```
2. Preferred
```

```
3. Required
```

```
[1]> 3
```

```
You have chosen to enable TLS. Please use the 'certconfig' command to ensure that there is a valid certificate configured.
```

Note that this example asks you to use the `certconfig` command to ensure that there is a valid certificate that can be used with the listener. If you have not created any certificates, the listener uses the demonstration certificate that is pre-installed on the appliance. You may enable TLS with the demonstration certificate for testing purposes, but it is not secure and is not recommended for general use. Use the `listenerconfig -> edit -> certificate` command to assign a certificate to the listener.

Once you have configured TLS, the setting will be reflected in the summary of the listener in the CLI:

```
Name: Inboundmail
```

```
Type: Public
```

```
Interface: PublicNet (192.168.2.1/24) TCP Port 25
```

```
Protocol: SMTP
```

```
Default Domain:
```

```
Max Concurrency: 1000 (TCP Queue: 50)
```

```
Domain map: disabled
```

```
TLS: Required
```

- Step 4** Issue the `commit` command to enable the change.
-

# Enabling TLS and Certificate Verification on Delivery

You can require that TLS is enabled for email delivery to specific domains using the Destination Controls page or the `destconfig` command.

In addition to TLS, you can require that the domain's server certificate is verified. This domain verification is based on a digital certificate used to establish the domain's credentials. The validation process involves two validation requirements:

- The chain of issuer certificates for the SMTP session ends in a certificate issued by a trusted certificate authority (CA).
- The Common Name (CN) listed on the certificate matches either the receiving machine's DNS name or the message's destination domain.

- or -

The message's destination domain matches one of the DNS names in the certificate's Subject Alternative Name (subjectAltName) extension, as described in RFC 2459. The matching supports wildcards as described in section 3.1 of RFC 2818.

A trusted CA is a third-party organization or company that issues digital certificates used to verify identity and distributes public keys. This provides an additional level of assurance that the certificate is issued by a valid and trusted identity.

You can configure your Email Security appliance to send messages to a domain over a TLS connection as an alternative to envelope encryption. See the “Cisco Email Encryption” chapter for more information.

You can specify a certificate for the appliance to use for all outgoing TLS connections. To specify the certificate, click **Edit Global Settings** on the Destination Controls page or use `destconfig -> setup` in the CLI. The certificate is a global setting, not a per-domain setting.

You can specify 5 different settings for TLS for a given domain when you include a domain using the Destination Controls page or the `destconfig` command. In addition to specifying whether exchanges with a domain are required or preferred to be TLS encoded, you can dictate whether validation of the domain is necessary. See [Table 23-3](#) for an explanation of the settings.

**Table 23-3**      **TLS Settings for Delivery**

| TLS Setting         | Meaning                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>      | The default TLS setting set using the Destination Controls page or the <code>destconfig -&gt; default</code> subcommand used for outgoing connections from the listener to the MTA for the domain.<br><br>The value “Default” is set if you answer “no” to the question: “Do you wish to apply a specific TLS setting for this domain?”                                                                                                                 |
| <b>1. No</b>        | TLS is not negotiated for outgoing connections from the interface to the MTA for the domain.                                                                                                                                                                                                                                                                                                                                                            |
| <b>2. Preferred</b> | TLS is negotiated from the Email Security appliance interface to the MTA(s) for the domain. However, if the TLS negotiation fails (prior to receiving a 220 response), the SMTP transaction will continue “in the clear” (not encrypted). No attempt is made to verify if the certificate originates from a trusted certificate authority. If an error occurs after the 220 response is received the SMTP transaction does not fall back to clear text. |

**Table 23-3**      **TLS Settings for Delivery**

| <b>TLS Setting</b>           | <b>Meaning</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>3. Required</b>           | TLS is negotiated from the Email Security appliance interface to MTA(s) for the domain. No attempt is made to verify the domain's certificate. If the negotiation fails, no email is sent through the connection. If the negotiation succeeds, the mail is delivered via an encrypted session.                                                                                                                                                                                                                                                                                                          |
| <b>4. Preferred (Verify)</b> | <p>TLS is negotiated from the Email Security appliance to the MTA(s) for the domain. The appliance attempts to verify the domain's certificate.</p> <p>Three outcomes are possible:</p> <ul style="list-style-type: none"> <li>• TLS is negotiated and the certificate is verified. The mail is delivered via an encrypted session.</li> <li>• TLS is negotiated, but the certificate is not verified. The mail is delivered via an encrypted session.</li> <li>• No TLS connection is made and, subsequently the certificate is not verified. The email message is delivered in plain text.</li> </ul> |
| <b>5. Required (Verify)</b>  | <p>TLS is negotiated from the Email Security appliance to the MTA(s) for the domain. Verification of the domain's certificate is required.</p> <p>Three outcomes are possible:</p> <ul style="list-style-type: none"> <li>• A TLS connection is negotiated and the certificate is verified. The email message is delivered via an encrypted session.</li> <li>• A TLS connection is negotiated but the certificate is not verified by a trusted CA. The mail is not delivered.</li> <li>• A TLS connection is not negotiated. The mail is not delivered.</li> </ul>                                     |

If there is no specific entry for a given recipient domain in the good neighbor table, or if there is a specific entry but there is no specific TLS setting for the entry, then the behavior is whatever is set using the Destination Controls page or the `destconfig -> default` subcommand ("No," "Preferred," "Required," "Preferred (Verify)," or "Required (Verify)").

#### Related Topics

- [Sending Alerts When a Required TLS Connection Fails, page 23-11](#)
- [Logging, page 23-12](#)
- [CLI Example, page 23-12](#)

## Sending Alerts When a Required TLS Connection Fails

You can specify whether the Email Security appliance sends an alert if the TLS negotiation fails when delivering messages to a domain that requires a TLS connection. The alert message contains name of the destination domain for the failed TLS negotiation. The Email Security appliance sends the alert message to all recipients set to receive Warning severity level alerts for System alert types. You can manage alert recipients via the System Administration > Alerts page in the GUI (or via the `alertconfig` command in the CLI).



**Related Topics**

- [Enabling TLS Connection Alerts Using the GUI, page 23-12](#)
- [Enabling TLS Connection Alerts Using the CLI, page 23-12](#)

## Enabling TLS Connection Alerts Using the GUI

**Procedure**

- 
- |               |                                                                                                                                                                                                                                                                             |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Navigate to the Mail Policies Destination Controls page.                                                                                                                                                                                                                    |
| <b>Step 2</b> | Click <b>Edit Global Settings</b> .                                                                                                                                                                                                                                         |
| <b>Step 3</b> | Click <b>Enable</b> for “Send an alert when a required TLS connection fails.”<br><br>This is a global setting, not a per-domain setting. For information on the messages that the appliance attempted to deliver, use the Monitor > Message Tracking page or the mail logs. |
| <b>Step 4</b> | Submit and commit your changes.                                                                                                                                                                                                                                             |
- 

## Enabling TLS Connection Alerts Using the CLI

To enable TLS connection alerts using the CLI, use the `destconfig -> setup` command.

## Logging

The Email Security appliance will note in the mail logs instances when TLS is required for a domain but could not be used. Information on why the TLS connection could not be used will be included. The mail logs will be updated when any of the following conditions are met:

- The remote MTA does not support ESMTP (for example, it did not understand the EHLO command from the Email Security appliance).
- The remote MTA supports ESMTP but “STARTTLS” was not in the list of extensions it advertised in its EHLO response.
- The remote MTA advertised the “STARTTLS” extension but responded with an error when the Email Security appliance sent the STARTTLS command.

## CLI Example

In this example, the `destconfig` command is used to require TLS connections and encrypted conversations for the domain “partner.com.” The list is then printed.

A certificate for example.com is used for outgoing TLS connections instead of the demonstration certificate that is pre-installed. You may enable TLS with the demonstration certificate for testing purposes, but it is not secure and is not recommended for general use.

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

There is currently 1 entry configured.

Choose the operation you want to perform:

- SETUP - Change global settings.
- NEW - Create a new entry.
- DELETE - Remove an entry.
- DEFAULT - Change the default.
- LIST - Display a summary list of all entries.
- DETAIL - Display details for one destination or all entries.
- IMPORT - Import tables from a file.
- EXPORT - Export tables to a file.

[>>] **setup**

The "Demo" certificate is currently configured. You may use "Demo", but this will not be secure.

1. example.com

2. Demo

Please choose the certificate to apply:

[1]> **1**

Do you want to send an alert when a required TLS connection fails? [N]>

There is currently 1 entry configured.

Choose the operation you want to perform:

- SETUP - Change global settings.
- NEW - Create a new entry.
- DELETE - Remove an entry.

- DEFAULT - Change the default.
- LIST - Display a summary list of all entries.
- DETAIL - Display details for one destination or all entries.
- IMPORT - Import tables from a file.
- EXPORT - Export tables to a file.

[> **new**

Enter the domain you wish to limit.

[> **partner.com**

Do you wish to configure a concurrency limit for partner.com? [Y]> **n**

Do you wish to apply a messages-per-connection limit to this domain? [N]> **n**

Do you wish to apply a recipient limit to this domain? [N]> **n**

Do you wish to apply a specific bounce profile to this domain? [N]> **n**

Do you wish to apply a specific TLS setting for this domain? [N]> **y**

Do you want to use TLS support?

1. No
2. Preferred
3. Required
4. Preferred (Verify)
5. Required (Verify)

[1]> **3**

You have chosen to enable TLS. Please use the 'certconfig' command to ensure that there is a valid certificate configured.

Do you wish to apply a specific bounce verification address tagging setting for this domain? [N]> **n**

Do you wish to apply a specific bounce profile to this domain? [N]> **n**

There are currently 2 entries configured.

Choose the operation you want to perform:

- SETUP - Change global settings.
- NEW - Create a new entry.
- EDIT - Modify an entry.
- DELETE - Remove an entry.
- DEFAULT - Change the default.
- LIST - Display a summary list of all entries.
- DETAIL - Display details for one destination or all entries.
- CLEAR - Remove all entries.
- IMPORT - Import tables from a file.
- EXPORT - Export tables to a file.

[> **list**

|             | Rate     |       | Bounce       | Bounce    |
|-------------|----------|-------|--------------|-----------|
| Domain      | Limiting | TLS   | Verification | Profile   |
| =====       | =====    | ===== | =====        | =====     |
| partner.com | Default  | Req   | Default      | Default   |
| (Default)   | On       | Off   | Off          | (Default) |

There are currently 2 entries configured.

Choose the operation you want to perform:

- SETUP - Change global settings.
  - NEW - Create a new entry.
  - EDIT - Modify an entry.
  - DELETE - Remove an entry.
  - DEFAULT - Change the default.
  - LIST - Display a summary list of all entries.
  - DETAIL - Display details for one destination or all entries.
  - CLEAR - Remove all entries.
  - IMPORT - Import tables from a file.
  - EXPORT - Export tables to a file.
- [ ]>

## Managing Lists of Certificate Authorities

The appliance uses stored trusted certificate authorities that it uses to verify a certificate from a remote domain to establish the domain's credentials. You can configure the appliance to use the following trusted certificate authorities:

- **Pre-installed list.** The appliance has a pre-installed list of trusted certificate authorities. This is called the system list.
- **User-defined list.** You can customize a list of trusted certificate authorities and then import the list onto the appliance.

You can use either the system list or the customized list, and you can also use both lists to verify certificate from a remote domain.

Manage the lists using the Network > Certificates > Edit Certificate Authorities page in the GUI or the `certconfig > certauthority` command in the CLI.

On the Network > Certificates > Edit Certificate Authorities page, you can perform the following tasks:

- **View the system list (pre-installed) of certificate authorities.** For more information, see [Viewing the Pre-Installed list of Certificate Authorities, page 23-17](#).
- **Choose whether or not to use the system list.** You can enable or disable the system list. For more information, see [Disabling the System Certificate Authority List, page 23-17](#).
- **Choose whether or not to use a custom certificate authority list.** You can enable the appliance to use a custom list and then import the list from a text file. For more information, see [Importing a Custom Certificate Authority List, page 23-17](#).
- **Export the list of certificate authorities to a file.** You can export either the system or customized list of certificate authorities to a text file. For more information, see [Exporting a Certificate Authorities List, page 23-18](#).

**Related Topics**

- [Viewing the Pre-Installed list of Certificate Authorities, page 23-17](#)
- [Disabling the System Certificate Authority List, page 23-17](#)
- [Importing a Custom Certificate Authority List, page 23-17](#)
- [Exporting a Certificate Authorities List, page 23-18](#)

## Viewing the Pre-Installed list of Certificate Authorities

**Procedure**

- 
- |               |                                                                    |
|---------------|--------------------------------------------------------------------|
| <b>Step 1</b> | Navigate to the Network > Certificates page.                       |
| <b>Step 2</b> | Click <b>Edit Settings</b> in the Certificate Authorities section. |
| <b>Step 3</b> | Click <b>View System Certificate Authorities</b> .                 |
- 

## Disabling the System Certificate Authority List

The pre-installed system certificate authorities list cannot be removed from the appliance, but you can enable or disable it. You might want to disable it to allow the appliance to only use your custom list to verify certificates from remote hosts.

**Procedure**

- 
- |               |                                                                    |
|---------------|--------------------------------------------------------------------|
| <b>Step 1</b> | Navigate to the Network > Certificates page.                       |
| <b>Step 2</b> | Click <b>Edit Settings</b> in the Certificate Authorities section. |
| <b>Step 3</b> | Click <b>Disable</b> for the System List.                          |
| <b>Step 4</b> | Submit and commit your changes.                                    |
- 

## Importing a Custom Certificate Authority List

You can create a custom of list trusted certificate authorities and import it onto the appliance. The file must be in the PEM format and include certificates for the certificate authorities that you want the appliance to trust.

**Procedure**

- 
- |               |                                                                       |
|---------------|-----------------------------------------------------------------------|
| <b>Step 1</b> | Navigate to the Network > Certificates page.                          |
| <b>Step 2</b> | Click <b>Edit Settings</b> in the Certificate Authorities section.    |
| <b>Step 3</b> | Click <b>Enable</b> for the Custom List.                              |
| <b>Step 4</b> | Enter the full path to the custom list on a local or network machine. |

- Step 5** Submit and commit your changes.
- 

## Exporting a Certificate Authorities List

If you want to use only a subset of the trusted certificate authorities in the system or edit an existing custom list, you can export the list to a .txt file and edit it to add or remove certificate authorities. After you have finished editing the list, import the file back onto the appliance as a custom list.

### Procedure

---

- Step 1** Navigate to the Network > Certificates page.
- Step 2** Click **Edit Settings** in the Certificate Authorities section.
- Step 3** Click **Export List**.
- AsyncOS displays the Export Certificate Authority List page.
- Step 4** Select the list you want to export.
- Step 5** Enter a filename for the list.
- Step 6** Click **Export**.
- AsyncOS displays a dialog box asking if want to open or save the list as a .txt file.
- 

## Enabling a Certificate for HTTPS

You can enable a certificate for HTTPS services on an IP interface using either the Network > IP Interfaces page in the GUI or the `interfaceconfig` command in the CLI. When adding an IP interface via the GUI, select a certificate that you want to use for the HTTPS service, check the **HTTPS** check box, and enter the port number.

In following example, the `interfaceconfig` command is used to edit the IP interface PublicNet to enable HTTPS services on port 443 (the default port). All other defaults for the interface are accepted. (Typing Enter at the prompt accepts the default value shown in brackets.)

Note that this example shows using the demonstration certificate that is pre-installed on the appliance. You may enable HTTPS services with the demonstration certificate for testing purposes, but it is not secure and is not recommended for general use.



### Note

You can enable HTTPS services using the System Setup Wizard in the GUI. Refer to “Define the Default Router (Gateway), Configure the DNS Settings, and Enabling Secure Web Access” in the “Setup and Installation” chapter.

---

After the changes from this command are committed, users can access the Graphical User Interface (GUI) using the URL for secure HTTPS: `https://192.168.2.1`

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

Currently configured interfaces:

1. Management (192.168.42.42/24: mail3.example.com)
2. PrivateNet (192.168.1.1/24: mail3.example.com)
3. PublicNet (192.168.2.1/24: mail3.example.com)

Choose the operation you want to perform:

- NEW - Create a new interface.
- EDIT - Modify an interface.
- GROUPS - Define interface groups.
- DELETE - Remove an interface.

```
[> edit
```

Enter the number of the interface you wish to edit.

```
[> 3
```

IP interface name (Ex: "InternalNet"):

```
[PublicNet]>
```

Would you like to configure an IPv4 address for this interface (y/n)? [Y]> **y**

IPv4 Address (Ex: 192.168.1.2):

```
[192.168.2.1]>
```

Netmask (Ex: "255.255.255.0" or "0xffffffff"):

```
[24]>
```



```
Would you like to configure an IPv6 address for this interface (y/n)? [N]>
```

```
Ethernet interface:
```

1. Data 1
2. Data 2
3. Management

```
[2]>
```

```
Hostname:
```

```
[mail3.example.com]>
```

```
Do you want to enable Telnet on this interface? [N]>
```

```
Do you want to enable SSH on this interface? [N]>
```

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

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

```
Which port do you want to use for HTTP?
```

```
[80]>
```

```
Do you want to enable HTTPS on this interface? [N]> y
```

```
Which port do you want to use for HTTPS?
```

```
[443]> 443
```

```
Do you want to enable Spam Quarantine HTTP on this interface? [N]>
```

Do you want to enable Spam Quarantine HTTPS on this interface? [N]>

The "Demo" certificate is currently configured. You may use "Demo", but this will not be secure. To assure privacy, run "certconfig" first.

Both HTTP and HTTPS are enabled for this interface, should HTTP requests redirect to the secure service? [Y]>

Currently configured interfaces:

1. Management (192.168.42.42/24: mail3.example.com)
2. PrivateNet (192.168.1.1/24: mail3.example.com)
3. PublicNet (192.168.2.1/24: mail3.example.com)

Choose the operation you want to perform:

- NEW - Create a new interface.
- EDIT - Modify an interface.
- GROUPS - Define interface groups.
- DELETE - Remove an interface.

[ ]>





## Configuring Routing and Delivery Features

- [Routing Email for Local Domains, page 24-1](#)
- [Rewriting Addresses, page 24-6](#)
- [Creating Alias Tables, page 24-7](#)
- [Configuring Masquerading, page 24-16](#)
- [The Domain Map Feature, page 24-28](#)
- [Directing Bounced Email, page 24-35](#)
- [Controlling Email Delivery Using Destination Controls, page 24-42](#)
- [Bounce Verification, page 24-51](#)
- [Set Email Delivery Parameters, page 24-56](#)
- [Configuring Mail Gateways for all Hosted Domains Using Virtual Gateway™ Technology, page 24-59](#)
- [Using Global Unsubscribe, page 24-69](#)

### Routing Email for Local Domains

In [Chapter 5, “Configuring the Gateway to Receive Email”](#) you customized private and public listeners to service SMTP connections for an Enterprise Gateway configuration. Those listeners were customized to handle specific connections (via HAT modification) and receive mail for specific domains (via RAT modification of public listeners).

The appliance routes mail to local domains to hosts specified via the Network > SMTP Routes page (or the `smtproutes` command). This feature is similar to the `sendmail mailertable` feature.



#### Note

If you have completed the GUI's System Setup Wizard (or the Command Line Interface `systemsetup` command) as described in the “Setup and Installation” chapter and committed the changes, you defined the first SMTP route entries on the appliance for each RAT entry you entered at that time.

#### Related Topics

- [SMTP Routes Overview, page 24-2](#)
- [Default SMTP Route, page 24-2](#)
- [Defining an SMTP Route, page 24-3](#)

- [SMTP Routes Limits, page 24-3](#)
- [SMTP Routes and DNS, page 24-3](#)
- [SMTP Routes and Alerts, page 24-4](#)
- [SMTP Routes, Mail Delivery, and Message Splintering, page 24-4](#)
- [SMTP Routes and Outbound SMTP Authentication, page 24-4](#)
- [Managing SMTP Routes to Send Outbound Email Using the GUI, page 24-4](#)

## SMTP Routes Overview

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 AsyncOS operating system allows up to forty thousand (40,000) SMTP Route mappings to be configured for your appliance. (See [SMTP Routes Limits, page 24-3](#).)

This feature also allows host “globbing.” If you specify a partial domain, such as `.example.com`, then any domain ending in `example.com` matches the entry. For instance, `fred@foo.example.com` and `wilma@bar.example.com` both match the mapping.

If a host is not found in the SMTP Routes table, an MX lookup is performed using DNS. The result is not re-checked against the SMTP Routes table. If the DNS MX entry for `foo.domain` is `bar.domain`, any email sent to `foo.domain` is delivered to the host `bar.domain`. If you create a mapping for `bar.domain` to some other host, email addressed to `foo.domain` is not affected.

In other words, recursive entries are not followed. If there is an entry for `a.domain` to redirect to `b.domain`, and a subsequent entry to redirect email for `b.domain` to `a.domain`, a mail loop will *not* be created. In this case, email addressed to `a.domain` will be delivered to the MX host specified by `b.domain`, and conversely email addressed to `b.domain` will be delivered to the MX host specified by `a.domain`.

The SMTP Routes table is read from the top down for every email delivery. The most specific entry that matches a mapping wins. For example, if there are mappings for both `host1.example.com` and `.example.com` in the SMTP Routes table, the entry for `host1.example.com` will be used because it is the more specific entry — even if it appears after the less specific `.example.com` entry. Otherwise, the system performs a regular MX lookup on the domain of the Envelope Recipient.

## Default SMTP Route

You can also define a default SMTP route with the special keyword `ALL`. If a domain does not match a previous mapping in the SMTP Routes list, it defaults to being redirected to the MX host specified by the `ALL` entry.

When you print the SMTP Routes entries, the default SMTP route is listed as `ALL:`. You cannot delete the default SMTP route; you may only clear any values entered for it.

Configure the default SMTP route via the Network > SMTP Routes page or the `smtproutes` command.

## Defining an SMTP Route

Use the Network > SMTP Routes page (or the `smtproutes` command) to construct routes. When you create a new route, you first specify the domain or partial domain for which you want to create a permanent route. You then specify destination hosts. Destination hosts can be entered as fully-qualified hostnames or as IP addresses. IP addresses can be either Internet Protocol version 4 (IPv4) or version 6 (IPv6).

For IPv6 addresses, AsyncOS supports the following formats:

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

You can also specify a special destination host of `/dev/null` to drop the messages that match the entry. (So, in effect, specifying `/dev/null` for the default route will ensure that no mail received by the appliance is ever delivered.)

A receiving domain can have multiple destination hosts, each assigned a priority number, much like an MX record. The destination host with the lowest number identifies as the primary destination host for the receiving domain. Other destination hosts listed will be used as backup.

Destinations with identical priority will be used in a “round-robin” fashion. The round-robin process is based on SMTP connections, and is not necessarily message-based. Also, if one or more of the destination hosts are not responding, messages will be delivered to one of the reachable hosts. If all the configured destination hosts are not responding, mail is queued for the receiving domain and delivery to the destination hosts is attempted later. (It does not fail over to using MX records).

When constructing routes using the `smtproutes` command in the CLI, you can prioritize each destination host by using `/pri=`, followed by an integer between 0 and 65535 to assign priority (0 is the highest priority) after the hostname or IP address. For example, `host1.example.com/pri=0` has a higher priority than `host2.example.com/pri=10`. Separate multiple entries with commas.

## SMTP Routes Limits

You can define up to 40,000 routes. The final default route of `ALL` is counted as a route against this limit. Therefore, you can define up to 39,999 custom routes and one route that uses the special keyword `ALL`.

## SMTP Routes and DNS

Use the special keyword `USEDNS` to tell the appliance to do MX lookups to determine next hops for specific domains. This is useful when you need to route mail for subdomains to a specific host. For example, if mail to `example.com` is to be sent to the company’s Exchange server, you might have something similar to the following SMTP route:

```
example.com exchange.example.com
```

However, for mail to various subdomains (`foo.example.com`), add an SMTP route that looks like this:

```
.example.com USEDNS
```

## SMTP Routes and Alerts

Alerts sent from the appliance to addresses specified in the System Administration > Alerts page (or the `alertconfig` command) follow SMTP Routes defined for those destinations.

## SMTP Routes, Mail Delivery, and Message Splintering

Incoming: if one message has 10 recipients and they are all on the same Exchange server, AsyncOS will open one TCP connection and present exactly one message to the mail store, not 10 separate messages.

Outgoing: works similarly, but if one message is going to 10 recipients in 10 different domains, AsyncOS will open 10 connections to 10 MTAs and deliver them one email each.

Splintering: if one incoming message has 10 recipients and they are each in separate Incoming Policy groups (10 groups), the message will splinter even if all 10 recipients are on the same Exchange server. Thus, 10 separate emails will be delivered over a single TCP connection.

## SMTP Routes and Outbound SMTP Authentication

If an Outbound SMTP Authentication profile has been created, you can apply it to an SMTP Route. This allows authentication for outgoing mail in cases where the appliance sits behind a mail relay server that is at the edge of the network. For more information about Outbound SMTP Authentication, see [Outgoing SMTP Authentication, page 25-39](#).

## Managing SMTP Routes to Send Outbound Email Using the GUI

Use the Network > SMTP Routes page to manage SMTP Routes on your appliance. You can add, modify, and delete mappings in the table. You can export or import the SMTP Routes entries.

### Related Topics

- [Adding SMTP Routes, page 24-4](#)
- [Exporting SMTP Routes, page 24-5](#)
- [Importing SMTP Routes, page 24-5](#)

## Adding SMTP Routes

### Procedure

- |               |                                                                                                                                                                                                           |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Click <b>Add Route</b> on the Network > SMTP Routes page.                                                                                                                                                 |
| <b>Step 2</b> | Enter a receiving domain. This can be a hostname, domain, IPv4 address, or IPv6 address.                                                                                                                  |
| <b>Step 3</b> | Enter a destination host. This can be a hostname, IPv4 address, or IPv6 address. You can add multiple destination hosts by clicking <b>Add Row</b> and entering the next destination host in the new row. |



### Note

You can specify a port number by adding “:<port number>” to the destination host:  
`example.com:25`.

- Step 4** If you add multiple destination hosts, enter an integer between 0 and 65535 to assign priority to the hosts. 0 is the highest priority. See [Defining an SMTP Route, page 24-3](#) for more information.
- Step 5** Submit and commit your changes.
- 

## Exporting SMTP Routes

Similar to the Host Access Table (HAT) and the Recipient Access Table (RAT), you can also modify SMTP routes mappings by exporting and importing a file. To export the SMTP Routes:

### Procedure

---

- Step 1** Click **Export SMTP Routes** on the SMTP Routes page.
- Step 2** Enter a name for the file and click **Submit**.
- 

## Importing SMTP Routes

Similar to the Host Access Table (HAT) and the Recipient Access Table (RAT), you can also modify SMTP routes mappings by exporting and importing a file. To import SMTP Routes:

### Procedure

---

- Step 1** Click **Import SMTP Routes** on the SMTP Routes page.
- Step 2** Select the file that contains the exported SMTP Routes.
- Step 3** Click **Submit**. You are warned that importing will replace all existing SMTP Routes. All of the SMTP Routes in the text file are imported.
- Step 4** Click **Import**.

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


```
this is a comment, but the next line is not
```

```
ALL:
```

---

At this point, our Email Gateway configuration looks like this:



**Figure 24-1 SMTP Routes Defined for a Public Listener**

## Rewriting Addresses

AsyncOS provides several methods for rewriting Envelope Sender and Recipient addresses in the email pipeline. Rewriting addresses can be used, for example, to redirect mail sent to a partner domain or to hide (“mask”) your internal infrastructure.

Table 24-1 provides an overview of the various features used for rewriting sender and recipient email addresses.

**Table 24-1**      **Methods for Rewriting Addresses**

| Original Address | Change to   | Feature                                                                  | Works on                                                                                                                                                       |
|------------------|-------------|--------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| *@anydomain      | user@domain | Alias Tables (see <a href="#">Creating Alias Tables, page 24-7</a> )     | <ul style="list-style-type: none"> <li>Envelope Recipients only</li> <li>Applied globally</li> <li>Maps aliases to email addresses or other aliases</li> </ul> |
| *@olddomain      | *@newdomain | Domain Mapping (see <a href="#">The Domain Map Feature, page 24-28</a> ) | <ul style="list-style-type: none"> <li>Envelope Recipients only</li> <li>Applied per listener</li> </ul>                                                       |
| *@olddomain      | *@newdomain | Masquerading (see <a href="#">Configuring Masquerading, page 24-16</a> ) | <ul style="list-style-type: none"> <li>Envelope Sender and the To:, From:, and/or CC: headers</li> <li>Applied per listener</li> </ul>                         |

## Creating Alias Tables

Alias tables provide a mechanism to redirect messages to one or more recipients. You can construct a mapping table of aliases to usernames and other aliases in a similar fashion to the `/etc/mail/aliases` feature of a sendmail configuration on some Unix systems.

When the Envelope Recipient (also known as the Envelope To, or `RCPT TO`) of an email accepted by a listener matches an alias as defined in an alias table, the Envelope Recipient address of the email will be rewritten.



### Note

A listener checks the alias table and modifies the recipients *after* checking the RAT and *before* message filters. See the “Understanding the Email Pipeline” chapter.



### Note

The Alias Table functionality actually rewrites the Envelope Recipient of the email. This is different than the `smtproutes` command (see [Directing Bounced Email, page 24-35](#)), which does not rewrite the Envelope Recipient of the email, but instead simply reroutes the email to specified domains.

### Related Topics

- [Configuring an Alias Table from the Command Line, page 24-7](#)
- [Exporting and Importing an Alias Table, page 24-8](#)
- [Deleting Entries from the Alias Table, page 24-9](#)

## Configuring an Alias Table from the Command Line

Alias tables are defined in sections as follows: each section is headed by a domain context, which is a list of domains that the section is relevant to, followed by a list of maps.

A domain context is a list of one or more domains or partial domains, separated by commas and enclosed in square brackets ('[' and ']'). A domain is a string containing letters, digits hyphens, and periods as defined in RFC 1035, section 2.3.1., “Preferred name syntax.” A partial domain, such as `.example.com` is a domain that begins with a period. All domains that end with a substring matching the partial domain are considered a match. For example, the domain context `.example.com` would match `mars.example.com` and `venus.example.com`. Below the domain context is a list of maps, which are aliases followed by a list of recipients. A map is constructed as follows:

**Table 24-2**      **Alias Table Syntax**

| Left-hand Side (LHS)                   | Separator                    | Right-hand Side (RHS)                                   |
|----------------------------------------|------------------------------|---------------------------------------------------------|
| a list of one or more aliases to match | the colon character<br>(':') | a list of one or more recipient<br>addresses or aliases |

An alias in the **left-hand side** can contain the following formats:

|                    |                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>username</b>    | Specifies an alias to match. There must be a preceding “domains” attribute specified in the table. The lack of this parameter will produce an error. |
| <b>user@domain</b> | Specifies an exact email address to match on.                                                                                                        |

You can enter multiple aliases, separated by commas on a single left-hand side line.

Each recipient in the **right-hand side** can be a full `user@domain` email address, or another alias.

An alias file can contain “global” aliases (aliases that are applied globally instead of to a specific domain) with no implied domain, domain contexts within which aliases have one or more implied domains, or both.

“Chains” (or recursive entries) of aliases may be created, but they must end in a full email address.

A special destination of `/dev/null` is supported to drop the message in order to be compatible with context of a sendmail configuration. If a message is mapped to `/dev/null` via an alias table, the dropped counter is increased. (See the “Managing and Monitoring via the CLI” chapter.) The recipient is accepted but not enqueued.

#### Related Topics

- [Example Alias Table, page 24-9](#)
- [Example aliasconfig Command, page 24-11](#)

## Exporting and Importing an Alias Table

To import an alias table, first see [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) to ensure that you can access the appliance.

Use the `export` subcommand of the `aliasconfig` command to save any existing alias table. A file (whose name you specify) will be written to the `/configuration` directory for the listener. You can modify this file outside of the CLI and then re-import it. (If you have malformed entries in the file, errors are printed when you try to import the file.)

Place the alias table file in the `/configuration` directory, and then use the `import` subcommand of the `aliasconfig` command to upload the file.

Comment out lines in the table using a number symbol (#) at the beginning of each line.

Remember to issue the `commit` command after you import an alias table file so that the configuration changes take effect.

## Deleting Entries from the Alias Table

If you delete entries from the alias table from the command line interface (CLI), you are prompted to choose a domain group first. Choose the “ALL (any domain)” entry to see a numbered list of aliases that apply to all domains. Then choose the number(s) of the aliases you want to delete.

### Example Alias Table

**Note**

All entries in this example table have been commented out.

```
sample Alias Table file

copyright (c) 2001-2005, IronPort Systems, Inc.

#

Incoming Envelope To addresses are evaluated against each
entry in this file from top to bottom. The first entry that
matches will be used, and the Envelope To will be rewritten.

#

Separate multiple entries with commas.

#

Global aliases should appear before the first domain
context. For example:

#

admin@example.com: administrator@example.com
postmaster@example.net: administrator@example.net
#

This alias has no implied domain because it appears
before a domain context:

#

someaddr@somewhere.dom: specificperson@here.dom
#
```

```
The following aliases apply to recipients @ironport.com and
any subdomain within .example.com because the domain context
is specified.
#
Email to joe@ironport.com or joe@foo.example.com will
be delivered to joseph@example.com.
#
Similarly, email to fred@mx.example.com will be
delivered to joseph@example.com
#
[ironport.com, .example.com]
#
joe, fred: joseph@example.com
#

In this example, email to partygoers will be sent to
three addresses:
#
partygoers: wilma@example.com, fred@example.com, barney@example.com
#
In this example, mail to help@example.com will be delivered to
customercare@otherhost.dom. Note that mail to help@ironport.com will
NOT be processed by the alias table because the domain context
overrides the previous domain context.
#
[example.com]
#
help: customercare@otherhost.dom
#
```

```
In this example, mail to nobody@example.com is dropped.

#

nobody@example.com: /dev/null

#

"Chains" may be created, but they must end in an email address.

For example, email to "all" will be sent to 9 addresses:

#

[example.com]

#

all: sales, marketing, engineering

sales: joe@example.com, fred@example.com, mary@example.com

marketing:bob@example.com, advertising

engineering:betty@example.com, miles@example.com, chris@example.com

advertising:richard@example.com, karen@advertising.com
```

## Example aliasconfig Command

In this example, the `aliasconfig` command is used to construct an alias table. First, the domain context of **example.com** is specified. Then, an alias of **customercare** is constructed so that any email sent to `customercare@example.com` is redirected to `bob@example.com`, `frank@example.com`, and `sally@example.com`. Next, a global alias of **admin** is constructed so that an email sent to `admin` is redirected to `administrator@example.com`. Finally, the alias table is printed to confirm.

Note that when the table is printed, the global alias for `admin` appears *before* the first domain context of `example.com`.

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

```
No aliases in table.
```

```
Choose the operation you want to perform:
```

- NEW - Create a new entry.
- IMPORT - Import aliases from a file.

```
[> new
```

How do you want your aliases to apply?

1. Globally
2. Add a new domain context

[1]> **2**

Enter new domain context.

Separate multiple domains with commas.

Partial domains such as .example.com are allowed.

[> **example.com**

Enter the alias(es) to match on.

Separate multiple aliases with commas.

Allowed aliases:

- "user" - This user in this domain context.
- "user@domain" - This email address.

[> **customercare**

Enter address(es) for "customercare".

Separate multiple addresses with commas.

[> **bob@example.com, frank@example.com, sally@example.com**

Adding alias customercare: bob@example.com,frank@example.com,sally@example.com

Do you want to add another alias? [N]> **n**

There are currently 1 mappings defined.

Choose the operation you want to perform:

- NEW - Create a new entry.

- EDIT - Modify an entry.
- DELETE - Remove an entry.
- PRINT - Display the table.
- IMPORT - Import aliases from a file.
- EXPORT - Export table to a file.
- CLEAR - Clear the table.

[> **new**

How do you want your aliases to apply?

1. Globally
2. Add a new domain context
3. example.com

[1]> **1**

Enter the alias(es) to match on.

Separate multiple aliases with commas.

Allowed aliases:

- "user@domain" - This email address.
- "user" - This user for any domain
- "@domain" - All users in this domain.
- "@.partialdomain" - All users in this domain, or any of its sub domains.

[> **admin**

Enter address(es) for "admin".

Separate multiple addresses with commas.

[> **administrator@example.com**

Adding alias admin: administrator@example.com



```
Do you want to add another alias? [N]> n
```

```
There are currently 2 mappings defined.
```

```
Choose the operation you want to perform:
```

- NEW - Create a new entry.
- EDIT - Modify an entry.
- DELETE - Remove an entry.
- PRINT - Display the table.
- IMPORT - Import aliases from a file.
- EXPORT - Export table to a file.
- CLEAR - Clear the table.

```
[]> print
```

```
admin: administrator@example.com
```

```
[example.com]
```

```
customer: bob@example.com, frank@example.com, sally@example.com
```

```
There are currently 2 mappings defined.
```


```
Choose the operation you want to perform:
```

- NEW - Create a new entry.
- EDIT - Modify an entry.
- DELETE - Remove an entry.
- PRINT - Display the table.
- IMPORT - Import aliases from a file.
- EXPORT - Export table to a file.

- CLEAR - Clear the table.

[>

At this point, our Email Gateway configuration looks like this:

**Figure 24-2 Alias Tables Defined for the Appliance**

## Configuring Masquerading

Masquerading is a feature that rewrites the Envelope Sender (also known as the sender, or MAIL FROM) and the To:, From:, and/or CC: headers on email processed by a listener according to a table that you construct. A typical example implementation of this feature is “Virtual Domains,” which allows you to host multiple domains from a single site. Another typical implementation is “hiding” your network infrastructure by “stripping” the subdomains from strings in email headers. The Masquerading feature is available for both private and public listeners.

**Note**

The Masquerading feature is configured on a per-listener basis, as opposed to the Alias Tables functionality, which is configured for the entire system.

**Note**

A listener checks the masquerading table for matches and modifies the recipients while the message is in the work queue, immediately after LDAP recipient acceptance queries and before LDAP routing queries. See the “Understanding the Email Pipeline” chapter.

The Masquerading feature actually rewrites addresses for the Envelope Sender and the To:, From:, and CC: fields of the email that has been received. You can specify different masquerading parameters for each listener you create in one of two ways:

- via a static table of mappings you create
- via an LDAP query.

This section discusses the static table method. The table format is forward-compatible with the `/etc/mail/genericstable` feature of a sendmail configuration on some Unix systems. See [Chapter 25, “LDAP Queries”](#) for more information on LDAP masquerading queries.

**Related Topics**

- [Masquerading and altsrghost, page 24-17](#)

## Masquerading and altsrghost

Generally, the masquerading feature rewrites the Envelope Sender, and any subsequent actions to be performed on the message will be “triggered” from the masqueraded address. However, when you run the `altsrghost` command from the CLI, the `altsrghost` mappings are triggered from the original address (and not the modified, masqueraded address).

For more information, see [Configuring Mail Gateways for all Hosted Domains Using Virtual Gateway™ Technology, page 24-59](#) and [Review: Email Pipeline, page 24-74](#).

**Related Topics**

- [Configuring Static Masquerading Tables, page 24-17](#)
- [Sample Masquerading Table for a Private Listener, page 24-19](#)
- [Importing a Masquerading Table, page 24-19](#)
- [Example Masquerading, page 24-19](#)

## Configuring Static Masquerading Tables

You configure the static masquerading table of mappings by using the `edit -> masquerade` subcommand of the `listenerconfig` command. Alternatively, you can import a file containing the mappings. See [Importing a Masquerading Table, page 24-19](#). The subcommand creates and maintains a table that maps input addresses, usernames, and domains to new addresses and domains. See [Chapter 25, “LDAP Queries”](#) for more information on LDAP masquerading queries.

When messages are injected into the system, the table is consulted, and the message is rewritten if a match in the header is found.

A domain masquerading table is constructed as follows:

**Table 24-3 Masquerading Table Syntax**

| Left-hand Side (LHS)                                    | Separator                           | Right-hand Side (RHS)                |
|---------------------------------------------------------|-------------------------------------|--------------------------------------|
| a list of one or more usernames and/or domains to match | whitespace (space or tab character) | the rewritten username and/or domain |

The following table lists valid entries in the masquerading table:

| Left-hand Side (LHS)                                                                                                                                                                                                       | Right-hand Side (RHS) |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| <b>username</b>                                                                                                                                                                                                            | username@domain       |
| This entry specifies a username to match. Incoming email messages matching a username on the left-hand side are matched and rewritten with the address on the right-hand side. The right-hand side must be a full address. |                       |
| <b>user@domain</b>                                                                                                                                                                                                         | username@domain       |
| The entry specifies an exact address to match. Incoming messages matching a full address on the left-hand side are rewritten with the address listed on the right-hand side. The right-hand side must be a full address.   |                       |
| <b>@domain</b>                                                                                                                                                                                                             | @domain               |
| This entry specifies any address with the specified domain. The original domain on the left-hand side is replaced with the domain in the right-hand side, leaving the username intact.                                     |                       |
| <b>@.partialdomain</b>                                                                                                                                                                                                     | @domain               |
| This entry specifies any address with the specified domain. The original domain on the left-hand side is replaced with the domain in the right-hand side, leaving the username intact.                                     |                       |
| <b>ALL</b>                                                                                                                                                                                                                 | @domain               |

The **ALL** entry matches bare addresses and rewrites them with the address on the right-hand side. The right-hand side must be a domain preceded by an “@”. This entry always has the lowest precedence regardless of its location in the table.

**Note** You can use the **ALL** entry for private listeners only.

- Rules are matched by the order in which they appear in the masquerading table.
- Addresses in the From:, To:, and CC: fields in the headers are matched and rewritten upon receiving by default. You can also configure the option to match and rewrite the Envelope Sender. Enable and disable the Envelope Sender and which headers to rewrite using the `config` subcommand.
- You can comment out lines in the table using a number symbol (#) at the beginning of each line. Everything following a # to the end of the line will be considered a comment and ignored.
- A masquerading table is limited to 400,000 entries, whether you create them via the `new` subcommand or import them from a file.

## Sample Masquerading Table for a Private Listener

```
sample Masquerading file

@example.com @example.com # Hides local subdomains in the header

sales sales_team@success.com

@techsupport tech_support@biggie.com

user@localdomain user@company.com

ALL @bigsender.com
```

## Importing a Masquerading Table

A traditional sendmail `/etc/mail/genericstable` file can be imported. To import a `genericstable` file, first see [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) to ensure that you can access the appliance.

Place the `genericstable` file in the configuration directory, and then use the `import` subcommand of the `masquerade` subcommand to upload the file. Use the commands in this order:

```
listenerconfig -> edit -> listener_number -> masquerade -> import
```

Alternatively, you can use the `export` subcommand to download the existing configuration. A file (whose name you specify) will be written to the configuration directory. You can modify this file outside of the CLI and then import it again.

When you use the `import` subcommand, ensure that the file contains only valid entries. If there is an invalid entry (for example, a left-hand side with no right-hand side), the CLI reports syntax errors when you import the file. If there is a syntax error during import, no mappings in the entire file are imported.

Remember to issue the `commit` command after you import a `genericstable` file so that the configuration changes for the listener take effect.

## Example Masquerading

In this example, the `masquerade` subcommand of `listenerconfig` is used to construct a domain masquerading table for the private listener named “OutboundMail” on the PrivateNet interface.

First, the option to use LDAP for masquerading is declined. (For information on configuring LDAP masquerading queries, see [Chapter 25, “LDAP Queries”](#) for more information on LDAP masquerading queries.)

Then, a partial domain notation of `@.example.com` is mapped to `@example.com` so that any email sent from any machine in the subdomain of `.example.com` will be mapped to `example.com`. Then, the username `joe` is mapped to the domain `joe@example.com`. The domain masquerading table is then printed to confirm both entries, and then exported to a file named `masquerade.txt`. The `config` subcommand is used to disable re-writing addresses in the CC: field, and finally, the changes are committed.

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

```
Currently configured listeners:
```

1. InboundMail (on PublicNet, 192.168.2.1) SMTP TCP Port 25 Public
2. OutboundMail (on PrivateNet, 192.168.1.1) SMTP TCP Port 25 Private

Choose the operation you want to perform:

- NEW - Create a new listener.
- EDIT - Modify a listener.
- DELETE - Remove a listener.
- SETUP - Change global settings.

[> **edit**

Enter the name or number of the listener you wish to edit.

[> **2**

Name: OutboundMail

Type: Private

Interface: PrivateNet (192.168.1.1/24) TCP Port 25

Protocol: SMTP

Default Domain:

Max Concurrency: 600 (TCP Queue: 50)

Domain Map: Disabled

TLS: No

SMTP Authentication: Disabled

Bounce Profile: Default

Footer: None

LDAP: Off

Choose the operation you want to perform:

- NAME - Change the name of the listener.
- INTERFACE - Change the interface.

- LIMITS - Change the injection limits.
- SETUP - Configure general options.
- HOSTACCESS - Modify the Host Access Table.
- BOUNCECONFIG - Choose the bounce profile to use for messages injected on this listener.
- MASQUERADE - Configure the Domain Masquerading Table.
- DOMAINMAP - Configure domain mappings.
- LDAPACCEPT - Configure an LDAP query to determine whether a recipient address should be accepted or bounced/dropped.
- LDAPROUTING - Configure an LDAP query to reroute messages.
- LDAPGROUP - Configure an LDAP query to determine whether a sender or recipient is in a specified group.
- SMTPAUTH - Configure an SMTP authentication.

[> **masquerade**

Do you want to use LDAP for masquerading? [N]> **n**

Domain Masquerading Table

There are currently 0 entries.

Masqueraded headers: To, From, Cc

Choose the operation you want to perform:

- NEW - Create a new entry.
- DELETE - Remove an entry.
- PRINT - Display all entries.
- IMPORT - Import all entries from a file.
- EXPORT - Export all entries to a file.
- CONFIG - Configure masqueraded headers.
- CLEAR - Remove all entries.



```
[> new
```

Enter the source address or domain to masquerade.

Username like "joe" are allowed.

Full addresses like "user@example.com" are allowed.

Full addresses with subdomain wildcards such as "username@.company.com" are allowed.

Domains like @example.com and @.example.com are allowed.

Hosts like @training and @.sales are allowed.

```
[> @.example.com
```

Enter the masqueraded address or domain.

Domains like @example.com are allowed.

Full addresses such as user@example.com are allowed.

```
[> @example.com
```

Entry mapping @.example.com to @example.com created.

Domain Masquerading Table

There are currently 1 entries.

Masqueraded headers: To, From, Cc

Choose the operation you want to perform:

- NEW - Create a new entry.
- DELETE - Remove an entry.
- PRINT - Display all entries.
- IMPORT - Import all entries from a file.
- EXPORT - Export all entries to a file.
- CONFIG - Configure masqueraded headers.
- CLEAR - Remove all entries.

```
[> new
```

Enter the source address or domain to masquerade.

Username like "joe" are allowed.

Full addresses like "user@example.com" are allowed.

Full addresses with subdomain wildcards such as "username@.company.com" are allowed.

Domains like @example.com and @.example.com are allowed.

Hosts like @training and @.sales are allowed.

```
[> joe
```

Enter the masqueraded address.

Only full addresses such as user@example.com are allowed.

```
[> joe@example.com
```

Entry mapping joe to joe@example.com created.

Domain Masquerading Table

There are currently 2 entries.

Masqueraded headers: To, From, Cc

Choose the operation you want to perform:

- NEW - Create a new entry.
- DELETE - Remove an entry.
- PRINT - Display all entries.
- IMPORT - Import all entries from a file.
- EXPORT - Export all entries to a file.
- CONFIG - Configure masqueraded headers.
- CLEAR - Remove all entries.

```
[> print
```

```
@.example.com @example.com
```

```
joe joe@example.com
```

Domain Masquerading Table

There are currently 2 entries.

Masqueraded headers: To, From, Cc

Choose the operation you want to perform:

- NEW - Create a new entry.
- DELETE - Remove an entry.
- PRINT - Display all entries.
- IMPORT - Import all entries from a file.
- EXPORT - Export all entries to a file.
- CONFIG - Configure masqueraded headers.
- CLEAR - Remove all entries.

```
[> export
```

Enter a name for the exported file:

```
[> masquerade.txt
```

Export completed.

Domain Masquerading Table

There are currently 2 entries.

Masqueraded headers: To, From, Cc

Choose the operation you want to perform:

- NEW - Create a new entry.
- DELETE - Remove an entry.
- PRINT - Display all entries.
- IMPORT - Import all entries from a file.
- EXPORT - Export all entries to a file.
- CONFIG - Configure masqueraded headers.
- CLEAR - Remove all entries.

[> **config**

Do you wish to masquerade Envelope Sender?

[N]> **y**

Do you wish to masquerade From headers?

[Y]> **y**

Do you wish to masquerade To headers?

[Y]> **y**

Do you wish to masquerade CC headers?

[Y]> **n**

Do you wish to masquerade Reply-To headers?

[Y]> **n**

Domain Masquerading Table

There are currently 2 entries.

- NEW - Create a new entry.

- DELETE - Remove an entry.
- PRINT - Display all entries.
- IMPORT - Import all entries from a file.
- EXPORT - Export all entries to a file.
- CONFIG - Configure masqueraded headers.
- CLEAR - Remove all entries.

[ ]>

Name: OutboundMail

Type: Private

Interface: PrivateNet (192.168.1.1/24) TCP Port 25

Protocol: SMTP

Default Domain:

Max Concurrency: 600 (TCP Queue: 50)

Domain Map: Disabled

TLS: No

SMTP Authentication: Disabled

Bounce Profile: Default

Footer: None

LDAP: Off

Choose the operation you want to perform:

- NAME - Change the name of the listener.
- INTERFACE - Change the interface.
- LIMITS - Change the injection limits.
- SETUP - Configure general options.
- HOSTACCESS - Modify the Host Access Table.
- BOUNCECONFIG - Choose the bounce profile to use for messages injected on this listener.

- MASQUERADE - Configure the Domain Masquerading Table.
  - DOMAINMAP - Configure domain mappings.
  - LDAPACCEPT - Configure an LDAP query to determine whether a recipient address should be accepted or bounced/dropped.
  - LDAPROUTING - Configure an LDAP query to reroute messages.
  - LDAPGROUP - Configure an LDAP query to determine whether a sender or recipient is in a specified group.
  - SMTPAUTH - Configure an SMTP authentication.
- [ ]>

Currently configured listeners:


1. InboundMail (on PublicNet, 192.168.2.1) SMTP TCP Port 25 Public
2. OutboundMail (on PrivateNet, 192.168.1.1) SMTP TCP Port 25 Private

Choose the operation you want to perform:

- NEW - Create a new listener.
- EDIT - Modify a listener.
- DELETE - Remove a listener.
- SETUP - Change global settings.

[ ]>

Our Enterprise Gateway configuration now looks like this:

**Figure 24-3** Masquerading Defined for a Private Listener

## The Domain Map Feature

You can configure a “domain map” for listeners. For each listener you configure, you can construct a domain map table which rewrites the Envelope Recipient for each recipient in a message that matches a domain in the domain map table. This feature is similar to the sendmail “Domain Table” or Postfix “Virtual Table” feature. Only the Envelope Recipient is affected; the “To:” headers are not re-written by this feature.

**Note**

The processing of the domain map feature happens immediately before the RAT and right after Default Domain is evaluated. See the “Understanding the Email Pipeline” chapter.

A common implementation of the domain map feature is to accept incoming mail for more than one legacy domain. For example, if your company has acquired another company, you could construct a domain map on the appliance to accept messages for the acquired domain and rewrite the Envelope Recipients to your company’s current domain.

**Note**

You can configure up to 20,000 separate, unique domain mappings.

**Table 24-4 Domain Map Table Example Syntax**

| Left Side                         | Right Side                                                              | Comments                                         |
|-----------------------------------|-------------------------------------------------------------------------|--------------------------------------------------|
| <code>username@example.com</code> | <code>username2@example.net</code>                                      | Only complete address for the right side         |
| <code>user@.example.com</code>    | <code>user2@example.net</code>                                          |                                                  |
| <code>@example.com</code>         | <code>user@example.net</code><br><i>or</i><br><code>@example.net</code> | Complete address or fully-qualified domain name. |
| <code>@.example.com</code>        | <code>user@example.net</code><br><i>or</i><br><code>@example.net</code> |                                                  |

In the following example, the `domainmap` subcommand of the `listenerconfig` command is used to create a domain map for the public listener “InboundMail.” Mail for the domain and any subdomain of `oldcompanyname.com` is mapped to the domain `example.com`. The mapping is then printed for confirmation. Contrast this example with the configuration of placing both domains in the listener’s RAT: the domain map feature will actually rewrite the Envelope Recipient of `joe@oldcomapanyname.com` to `joe@example.com`, whereas placing the domain `oldcompanyname.com` in the listener’s RAT will simply accept the message for `joe@oldcompanyname.com` and route it without rewriting the Envelope Recipient. Also, contrast this example with the alias table feature. Alias tables *must* resolve to an explicit address; they cannot be constructed to map “any `username@domain`” to “the same `username@newdomain`.”

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

Currently configured listeners:

1. Inboundmail (on PublicNet, 192.168.2.1) SMTP TCP Port 25 Public
2. Outboundmail (on PrivateNet, 192.168.1.1) SMTP TCP Port 25 Private

Choose the operation you want to perform:



- NEW - Create a new listener.
- EDIT - Modify a listener.
- DELETE - Remove a listener.
- SETUP - Change global settings.

[> **edit**

Enter the name or number of the listener you wish to edit.

[> **1**

Name: InboundMail

Type: Public

Interface: PublicNet (192.168.2.1/24) TCP Port 25

Protocol: SMTP

Default Domain:

Max Concurrency: 1000 (TCP Queue: 50)

Domain Map: Disabled

TLS: No

SMTP Authentication: Disabled

Bounce Profile: Default

Use SenderBase For Reputation Filters and IP Profiling: Yes

Footer: None

LDAP: Off

Choose the operation you want to perform:

- NAME - Change the name of the listener.
- INTERFACE - Change the interface.
- LIMITS - Change the injection limits.
- SETUP - Configure general options.
- HOSTACCESS - Modify the Host Access Table.

- RCPTACCESS - Modify the Recipient Access Table.
- BOUNCECONFIG - Choose the bounce profile to use for messages injected on this listener.
- MASQUERADE - Configure the Domain Masquerading Table.
- DOMAINMAP - Configure domain mappings.

```
[> domainmap
```

Domain Map Table

There are currently 0 Domain Mappings.

Domain Mapping is: disabled

Choose the operation you want to perform:

- NEW - Create a new entry.
- IMPORT - Import domain mappings from a file.

```
[> new
```

Enter the original domain for this entry.

Domains such as "@example.com" are allowed.

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

Email addresses such as "test@example.com" and "test@.example.com"

are also allowed.

```
[> @.oldcompanyname.com
```

Enter the new domain for this entry.

The new domain may be a fully qualified

such as "@example.domain.com" or a complete

email address such as "test@example.com"

```
[> @example.com
```

```
Domain Map Table
```

```
There are currently 1 Domain Mappings.
```

```
Domain Mapping is: enabled
```

```
Choose the operation you want to perform:
```

- NEW - Create a new entry.
- EDIT - Modify an entry.
- DELETE - Remove an entry.
- PRINT - Display all domain mappings.
- IMPORT - Import domain mappings from a file.
- EXPORT - Export domain mappings to a file.
- CLEAR - Clear all domain mappings.

```
[> print
```

```
@.oldcompanyname.com --> @example.com
```

```
Domain Map Table
```

```
There are currently 1 Domain Mappings.
```

```
Domain Mapping is: enabled
```

```
Choose the operation you want to perform:
```

- NEW - Create a new entry.
- EDIT - Modify an entry.
- DELETE - Remove an entry.
- PRINT - Display all domain mappings.
- IMPORT - Import domain mappings from a file.

- EXPORT - Export domain mappings to a file.

- CLEAR - Clear all domain mappings.

[ ]>

Name: InboundMail

Type: Public

Interface: PublicNet (192.168.2.1/24) TCP Port 25

Protocol: SMTP

Default Domain:

Max Concurrency: 1000 (TCP Queue: 50)

Domain Map: Enabled

TLS: No

SMTP Authentication: Disabled

Bounce Profile: Default

Use SenderBase For Reputation Filters and IP Profiling: Yes

Footer: None

LDAP: Off

Choose the operation you want to perform:

- NAME - Change the name of the listener.

- INTERFACE - Change the interface.

- LIMITS - Change the injection limits.

- SETUP - Configure general options.

- HOSTACCESS - Modify the Host Access Table.

- RCPTACCESS - Modify the Recipient Access Table.

- BOUNCECONFIG - Choose the bounce profile to use for messages injected on this listener.

- MASQUERADE - Configure the Domain Masquerading Table.

- DOMAINMAP - Configure domain mappings.

[ ]>

#### Related Topics

- [Importing and Exporting a Domain Map Table, page 24-34](#)

## Importing and Exporting a Domain Map Table

To import or export a domain map table, first see [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) to ensure that you can access the appliance.

Create a text file of entries of domains to map. Separate the entries with white space (either a tab character or spaces). Comment out lines in the table using a number symbol (#) at the beginning of each line.

Place the file in the configuration directory, and then use the `import` subcommand of the `domain` subcommand to upload the file. Use the commands in this order:


```
listenerconfig -> edit -> inejtor_number -> domainmap -> import
```

Alternatively, you can use the `export` subcommand to download the existing configuration. A file (whose name you specify) will be written to the configuration directory. You can modify this file outside of the CLI and then import it again.

When you use the `import` subcommand, ensure that the file contains only valid entries. If there is an invalid entry (for example, a left-hand side with no right-hand side), the CLI reports syntax errors when you import the file. If there is a syntax error during import, no mappings in the entire file are imported.

Remember to issue the `commit` command after you import a domain map table file so that the configuration changes for the listener take effect.

Our Enterprise Gateway configuration now looks like this:

**Figure 24-4 Domain Map Defined for a Public Listener**

## Directing Bounced Email

Bounced email is an inevitable part of any email delivery. Your appliance is able to process bounced email in a number of highly configurable ways.

Please note, this section describes how to control how your appliance generates outgoing bounces (based on incoming mail). To control how your appliance controls incoming bounces (based on outgoing mail) use Bounce Verification (see [Bounce Verification](#), page 24-51).

**Related Topics**

- [Handling Undeliverable Email, page 24-36](#)
- [Creating a New Bounce Profile, page 24-40](#)
- [Applying Bounce Profiles to Listeners, page 24-41](#)

## Handling Undeliverable Email

The AsyncOS operating system classifies undeliverable email, or “bounced messages,” into the following categories:

**“Conversational” bounces:**

**The remote domain bounces the message during the initial SMTP conversation.**

|              |                                                                                                                                                                         |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Soft bounces | A message that is temporarily undeliverable. For example, a user’s mailbox may be full. These messages can be retried at a later time. (e.g. An SMTP 4XX error code.)   |
| Hard bounces | A message that is permanently undeliverable. For example, the user no longer exists for that domain. These messages will not be retried. (e.g. An SMTP 5XX error code.) |

**“Delayed” (or “Non-conversational”) bounces:**

**The remote domain accepts the message for delivery, only to bounce it at a later time.**

|              |                                                                                                                                                                         |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Soft bounces | A message that is temporarily undeliverable. For example, a user’s mailbox may be full. These messages can be retried at a later time. (e.g. An SMTP 4XX error code.)   |
| Hard bounces | A message that is permanently undeliverable. For example, the user no longer exists for that domain. These messages will not be retried. (e.g. An SMTP 5XX error code.) |

You use the Bounce Profiles page on the Network menu in the GUI (or the `bounceconfig` command) to configure how AsyncOS 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. (See [Chapter 9, “Using Message Filters to Enforce Email Policies”](#) for more information.)

**Related Topics**

- [Notes on Soft and Hard Bounces, page 24-36](#)
- [Bounce Profile Parameters, page 24-37](#)
- [Hard Bounces and the status Command, page 24-38](#)
- [Conversational Bounces and SMTP Routes Message Filter actions, page 24-38](#)
- [Example Bounce Profiles, page 24-39](#)
- [Delivery Status Notification Format, page 24-39](#)
- [Delay Warning Messages, page 24-40](#)
- [Delay Warning Messages and Hard Bounces, page 24-40](#)

## Notes on Soft and Hard Bounces

- For conversational soft bounces, a soft bounce event is defined as each time a recipient delivery temporarily fails. A single recipient may incur several soft bounce events. You use the Bounce Profiles page or the `bounceconfig` command to configure parameters for each soft bounce event. (See [Bounce Profile Parameters, page 24-37](#).)

- By default, the system generates a bounce message and sends it to the original sender for each hard bounced recipient. (The message is sent to the address defined in the Envelope Sender address of the message envelope. Envelope From is also commonly referred to as the Envelope Sender.) You can disable this feature and instead rely on log files for information about hard bounces. (See the “Logging” chapter.)
- Soft bounces become hard bounces after the maximum time in queue or the maximum number of retries, whichever comes first.

## Bounce Profile Parameters

When configuring a bounce profile, the following parameters control how conversational bounces are handled per message:

**Table 24-5**      **Bounce Profile Parameters**

|                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Maximum number of retries</b>                                   | The number of times the system should try to reconnect to the recipient host to re-deliver the soft bounced message before treating it as a hard bounced message. The default is 100 retries.                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Maximum number of seconds in queue</b>                          | The amount of time the system should spend trying connect to the recipient host to re-deliver the soft bounced message before treating it as a hard bounced message. The default is 259,200 seconds (72 hours).                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Initial number of seconds to wait before retrying a message</b> | The amount of time the system should wait before the first attempt to re-deliver the soft bounced message. The default is 60 seconds. Set the initial retry time to a high value to reduce the frequency of soft bounce attempts. Conversely, to increase the frequency, lower the value.                                                                                                                                                                                                                                                                                                                         |
| <b>Maximum number of seconds to wait before retrying a message</b> | The maximum amount of time the system should wait before trying to re-deliver the soft bounced message. The default is 3,600 seconds (1 hour). This is not the interval between each subsequent try; rather, it is another parameter that can be used to control the number of retries. The initial retry interval is limited on the high end by the maximum retry interval. If the calculated retry interval period exceeds the maximum retry interval then the maximum retry interval is used instead.                                                                                                          |
| <b>Hard bounce message generation format</b>                       | Specify whether hard bounce message generation is enabled or disabled. If it is enabled, you can choose the format of the message. By default, bounce messages generated use the DSN format (RFC 1894). You can select a custom notification template to use for bounce messages. For more information, see the “Text Resources” chapter.<br><br>You can also choose whether or not to parse the DSN status field from the bounce response. If you choose “Yes,” AsyncOS searches the bounce response for a DSN status code (RFC 3436) and uses the code in the Status field of the delivery status notification. |
| <b>Send delay warning messages</b>                                 | Specify whether or not to send delay warnings. If enabled, specify the minimum interval between messages as well as the maximum number of retries to send.<br><br>You can select a custom notification template to use for warning messages. For more information, see the “Text Resources” chapter.                                                                                                                                                                                                                                                                                                              |
| <b>Specify Recipient for Bounces</b>                               | You can bounce messages to an alternate address rather than the default of the Envelope Sender address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |



**Table 24-5 Bounce Profile Parameters (continued)**

|                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Use DomainKeys signing for bounce and delay messages</b>                                                                                                                                  | You can select a DomainKeys profile to use for signing bounce and delay messages. For information on DomainKeys, see <a href="#">DomainKeys and DKIM Authentication, page 20-1</a> .                                                                                                                                                                                                       |
| <b>Global Settings</b>                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Configure these settings via the Edit Global Settings link on the Bounce Profiles page or by editing the default bounce profile via the <code>bounceconfig</code> command in the CLI.</b> |                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Initial number of seconds to wait before retrying an unreachable host</b>                                                                                                                 | The amount of time the system should wait before retrying a host that is unreachable. The default is 60 seconds.                                                                                                                                                                                                                                                                           |
| <b>Max interval allowed between retries to an unreachable host</b>                                                                                                                           | The maximum amount of time the system should wait before retrying a host that is unreachable. The default is 3,600 seconds (1 hour). When the delivery initially fails due to the host being down, it will start with the minimum number of seconds retry value, and for each subsequent retry to the downed host, will increase the duration, up to this maximum number of seconds value. |

## Hard Bounces and the status Command

When hard bounce message generation is enabled, the following counters in the `status` and `status detail` commands increment each time the appliance generates a hard bounce message for delivery:

| Counters:              | Reset | Uptime | Lifetime |
|------------------------|-------|--------|----------|
| Receiving              |       |        |          |
| Messages Received      | 0     | 0      | 0        |
| Recipients Received    | 0     | 0      | 0        |
| Gen. Bounce Recipients | 0     | 0      | 0        |

For more information, see the “Monitoring and Managing via the CLI” chapter. When hard bounce message generation is disabled, none of these counters increments when a recipient hard bounces.



### Note

The Envelope Sender address of the message envelope is different than the From: in the message headers. AsyncOS can be configured to send hard bounce messages to an email address different than the Envelope Sender address.

## Conversational Bounces and SMTP Routes Message Filter actions

Mappings for SMTP Routes and message filter actions are not applied to the routing of SMTP bounce messages generated by the appliance as a result of a conversational bounce. When an appliance receives a conversational bounce message, it generates an SMTP bounce message back to the Envelope Sender of the original message. In this case, the appliance is actually generating the message, so any SMTP Routes that apply to an injected message for relaying do not apply.

## Example Bounce Profiles

Consider these two examples using different bounce profile parameters:

**Table 24-6**      **Example 1: Bounce Profile Parameters**

| Parameter                                     | Value                      |
|-----------------------------------------------|----------------------------|
| Max number of retries                         | 2                          |
| Max number of seconds in queue                | 259,200 seconds (72 hours) |
| Initial number of seconds before retrying     | 60 seconds                 |
| Max number of seconds to wait before retrying | 60 seconds                 |

In Example 1, the first recipient delivery attempt is made at  $t=0$ , immediately after the message is injected into the appliance. With the default initial retry time of 60 seconds, the first retry attempt is made approximately one minute later at  $t=60$ . The retry interval is calculated and it is determined to use the maximum retry interval of 60 seconds. Thus, the second retry attempt is made at approximately  $t=120$ . Immediately after this retry attempt, the system generates a hard bounce message for that recipient because the maximum number of retries is two.

**Table 24-7**      **Example 2: Bounce Profile Parameters**

| Parameter                                     | Value       |
|-----------------------------------------------|-------------|
| Max number of retries                         | 100         |
| Max number of seconds in queue                | 100 seconds |
| Initial number of seconds before retrying     | 60 seconds  |
| Max number of seconds to wait before retrying | 120 seconds |

In Example 2, the first delivery attempt is made at  $t=0$  and the first retry is made at  $t=60$ . The system hard bounces the message immediately before the next delivery attempt (scheduled to occur at  $t=120$ ) because it has exceeded the maximum time in queue of 100 seconds.

## Delivery Status Notification Format

Bounce messages generated by the system, by default, use the Delivery Status Notification (DSN) format for both hard and soft bounces. DSN is a format defined by RFC 1894 (see <http://www.faqs.org/rfcs/rfc1894.html>) that “defines a MIME content-type that may be used by a message transfer agent (MTA) or electronic mail gateway to report the result of an attempt to deliver a message to one or more recipients.” By default, the delivery status notification includes an explanation of the delivery status and the original message if the message size is less than 10k. If the message size is greater than 10k, the delivery status notification includes the message headers only. If the message headers exceed 10k, the delivery status notification truncates the headers. If you want include messages (or message headers) that are greater than 10k in the DSN, you can use the `max_bounce_copy` parameter in the `bounceconfig` command (this parameter is only available from the CLI).

## Delay Warning Messages

Time in Queue Messages (delay notification messages) generated by the system also use the DSN format. Change the default parameters by using the Bounce Profiles page on the Network menu (or the `bounceconfig` command) to edit existing or create new bounce profiles and change the default values for:

- The minimum interval between sending delay warning messages.
- The maximum number of delay warning messages to send per recipient.

## Delay Warning Messages and Hard Bounces

Note that it is possible to receive both a delay warning and a hard bounce for the same message *simultaneously*, if you have set a very small durations for both the “Maximum Time in Queue” setting and the minimum interval setting for “Send Delay Warning Messages.” Cisco Systems recommends using the default values for these settings as a minimum if you choose to enable sending of delay warning messages.

Further, delay warning messages and bounce messages originated by the appliance may be delayed by as much as 15 minutes during processing.

## Creating a New Bounce Profile

In the following example, a bounce profile named `bouncepr1` is created using the Bounce Profiles page. In this profile, all hard bounced messages are sent to the alternate address `bounce-mailbox@example.com`. Delay warnings messages are enabled. One warning message will be sent per recipient, and the default value of 4 hours (14400 seconds) between warning messages is accepted.

### Related Topics

- [Editing the Default Bounce Profile, page 24-40](#)
- [Example of a Minimalist Bounce Profile, page 24-40](#)

## Editing the Default Bounce Profile

You can edit any bounce profile by clicking its name in the Bounce Profiles listing. You can also edit the default bounce profile. In this example, the default profile is edited to increase the maximum number of seconds to wait before retrying unreachable hosts from 3600 (one hour) to 10800 (three hours):

## Example of a Minimalist Bounce Profile


In the following example, a bounce profile named `minimalist` is created. In this profile, messages are not retried when they bounce (zero maximum retries), and the maximum time to wait before retrying is specified. Hard bounce messages are disabled, and soft bounce warnings are not sent.

## Applying Bounce Profiles to Listeners

Once you have created a bounce profile, you can apply that profile to a listener using the Network > Listeners page or the `listenerconfig` command.

In the following example, the `bouncepr1` profile is applied to the `OutgoingMail` listener.

At this point, our Email Gateway configuration looks like this:

**Figure 24-5** Applying a Bounce Profile to a Private Listener

## Controlling Email Delivery Using Destination Controls

Uncontrolled high-volume email delivery can overwhelm recipient domains. AsyncOS gives you full control of message delivery by defining the number of connections your appliance will open or the number of messages your appliance will send to each destination domain.

Using the Destination Controls feature (Mail Policies > Destination Controls in the GUI, or the `destconfig` command in the CLI), you can control:

- [Rate Limiting, page 24-43](#)

- [TLS, page 24-43](#)
- [Bounce Verification, page 24-43](#)
- [Bounce Profile, page 24-43](#)

## Rate Limiting

- **Concurrent Connections:** number of simultaneous connections to remote hosts the appliance will attempt to open.
- **Maximum Messages Per Connection:** number of messages your appliance will send to a destination domain before the appliance initiates a new connection.
- **Recipients:** number of recipients the appliance will send to a given remote host in a given time period.
- **Limits:** how to apply the limits you have specified on a per-destination and per MGA hostname basis.

## TLS

- Whether TLS connections to remote hosts will be accepted, allowed, or required (see [Controlling TLS, page 24-46](#)).
- Whether to send an alert when TLS negotiation fails when delivering a message to a remote host that requires a TLS connection. This is a global setting, not a per-domain setting.
- Assign a TLS certificate to use for all outbound TLS connections to remote hosts.

## Bounce Verification

- Whether or not to perform address tagging via Bounce Verification (see [Bounce Verification, page 24-51](#)).

## Bounce Profile

- Which bounce profile should be used by the appliance for a given remote host (the default bounce profile is set via the Network > Bounce Profiles page).

You can also control the default settings for unspecified domains.

### Related Topics

- [Determining Which Interface is Used for Mail Delivery, page 24-43](#)
- [Default Delivery Limits, page 24-44](#)
- [Working with Destination Controls, page 24-44](#)

# Determining Which Interface is Used for Mail Delivery

Unless you specify the output interface via the `deliveryconfig` command or via a message filter (`alt-src-host`), or through the use of a virtual gateway, the output interface is selected by the AsyncOS routing table. Basically, selecting “auto” means to let AsyncOS decide.

In greater detail: local addresses are identified by applying the interface netmask to the interface IP address. Both of these are set via the Network > Interfaces page or by the `interfaceconfig` command (or during system setup). If the address space overlaps, the most specific netmask is used. If a destination is local, packets are sent via the appropriate local interface.

If the destination is not local, packets are sent to the default router (set via the Network > Routing page or with the `setgateway` command). The IP address of the default router is local. The output interface is determined by the rule for selecting the output interface for local addresses. For example, AsyncOS chooses the most specific IP address and netmask that include the default router's IP address.

The routing table is configured via the Network > Routing page (or via the `routeconfig` command). A matching entry in the routing table takes precedence over the default route. A more specific route takes precedence over a less specific route.

## Default Delivery Limits

Each outbound destination domain has its own outbound queue. Therefore, each domain has a separate set of concurrency limits as specified in the Destination Controls table. Further, each unique domain not listed specifically in the Destination Controls table uses another set of the “Default” limits as set in the table.

## Working with Destination Controls

Use the Mail Policies > Destination Controls page in the GUI or the `destconfig` command in the CLI to create, edit, and delete Destination Control entries.

### Related Topics

- [Controlling the Version of Internet Protocol Addresses, page 24-44](#)
- [Controlling the Number of Connections, Messages, and Recipients to a Domain, page 24-45](#)
- [Controlling TLS, page 24-46](#)
- [Controlling Bounce Verification Tagging, page 24-47](#)
- [Controlling Bounces, page 24-47](#)
- [Adding a New Destination Control Entry, page 24-47](#)
- [Importing and Exporting Destination Control Configurations, page 24-47](#)
- [Destination Controls and the CLI, page 24-51](#)

## Controlling the Version of Internet Protocol Addresses

You can configure which version of Internet Protocol addresses to use for the connection to a domain. The Email Security appliance uses both Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6). You can configure a listener on the appliance to use one version of the protocol or both.

If the “Required” setting for either IPv4 or IPv6 is specified, the appliance will negotiate a connection to the domain using an address of the specified version. If the domain doesn’t use that IP address version, no email will be sent. If the “Preferred” setting for either IPv4 or IPv6 is specified, the appliance will first attempt to negotiate a connection to the domain using an address of the specified version then fall back to the other if the first is not reachable.

## Controlling the Number of Connections, Messages, and Recipients to a Domain

You may want to limit how your appliance will deliver email to avoid overwhelming remote hosts or your own internal groupware servers with email from your appliance.

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 Controls or the `destconfig` command — previously the `setgoodtable` command). You can specify the domain name using the following syntax:

```
domain.com
```

or

```
.domain.com
```

This syntax enables AsyncOS to specify destination controls for sub-domains such as `sample.server.domain.com` without entering each full subdomain address individually.

For connections, messages, and recipients, you set whether the limits you define are enforced for each Virtual Gateway address, or for the entire system. (Virtual Gateway address limits control the number of concurrent connections per IP interface. System-wide limits control the total number of connections the appliance will allow.)

You also set whether the limits you define are enforced for each MX record of the specified domain, or for the entire domain. (Many domains have multiple MX records defined for accepting email.)



### Note

The current system default is 500 connections per domain and 50 messages per connection.

These values are explained in [Table 24-8](#).

**Table 24-8**      **Values in the Destination Controls Table**

| Field                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Concurrent Connections</b>          | The maximum number of outbound connections that will be made by the appliance to a given host. (Note that the domain can include your internal groupware hosts.)                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Maximum Messages Per Connection</b> | The maximum number of messages allowed for a single outbound connection from the appliance to a given host before initiating a new connection.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Recipients</b>                      | <p>The maximum number of recipients allowed within the given period of time. “None” denotes that there is no recipient limit for the given domain.</p> <p>The minimum period of time — between 1 and 60 minutes — that the appliance will count the number of recipients. Specifying a time period of “0” disables the feature.</p> <p><b>Note</b> If you change the recipient limit, AsyncOS resets the counters for all messages already in the queue. The appliance delivers the messages based on the new recipient limit.</p> |



**Table 24-8** *Values in the Destination Controls Table (continued)*

| Field               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Apply Limits</b> | <p>Specifies whether the limit will be applied (enforces) to the entire domain or to each mail exchange IP address specified for that domain. (Many domains have multiple MX records.)</p> <p>This setting applies to connection, message, and recipient limits.</p> <p>Specifies whether the limit will be applied system-wide or for each Virtual Gateway address.</p> <p><b>Note</b> If you have configured groups of IP addresses, but you have not configured virtual gateways, do not configure apply limits per each virtual gateway. This setting is intended only for systems configured to use virtual gateways. For information on configuring virtual gateways, see <a href="#">Configuring Mail Gateways for all Hosted Domains Using Virtual Gateway™ Technology</a>, page 24-59.</p> |

**Note**

If limits are applied per each Virtual Gateway address, you can still effectively implement system-wide limits by setting the Virtual Gateway limit to the system-wide limit you want divided by the number of possible virtual gateways. For example, if you have four Virtual Gateway addresses configured, and you do not want to open more than 100 simultaneous connections to the domain `yahoo.com`, set the Virtual Gateway limit to 25 simultaneous connections.

**Note**

The `delivernow` command, when acting on all domains, resets all counters tracked in the `destconfig` command.

## Controlling TLS

You can also configure the TLS (Transport Layer Security) on a per-domain basis. If the “Required” setting is specified, a TLS connection will be negotiated from the appliance listener to MTA(s) for the domain. If the negotiation fails, no email will be sent through the connection. For more information, see [Enabling TLS and Certificate Verification on Delivery](#), page 23-10.

You can specify whether the appliance sends an alert if the TLS negotiation fails when delivering messages to a domain that requires a TLS connection. The alert message contains name of the destination domain for the failed TLS negotiation. The appliance sends the alert message to all recipients set to receive Warning severity level alerts for System alert types. You can manage alert recipients via the System Administration > Alerts page in the GUI (or via the `alertconfig` command in the CLI).

To enable TLS connection alerts, click **Edit Global Settings** on the Destination Controls page or `destconfig -> setup` subcommand. This is a global setting, not a per-domain setting. For information on the messages that the appliance attempted to deliver, use the Monitor > Message Tracking page or the mail logs.

You must specify a certificate to use for all outgoing TLS connections. Use the **Edit Global Settings** on the Destination Controls page or `destconfig -> setup` subcommand to specify the certificate. For information on obtaining a certificate, see [Obtaining Certificates](#), page 23-2.

For more information on alerts, see the “System Administration” chapter.

## Controlling Bounce Verification Tagging

You can specify whether or not mail sent is tagged for bounce verification. You can specify this for the default, as well as specific destinations. Cisco suggests enabling bounce verification for the default, and then creating new destinations for specific exclusions. See [Bounce Verification, page 24-51](#) for more information.

## Controlling Bounces

In addition to controlling the number of connections and recipients will deliver to a remote host, you can also specify a bounce profile to be used for that domain. If specified, the bounce profile appears in the fifth column of the `destconfig` command. If you do not specify a bounce profile, the default bounce profile will be used. For more information, see [Creating a New Bounce Profile, page 24-40](#).

## Adding a New Destination Control Entry

### Procedure

- 
- |               |                                 |
|---------------|---------------------------------|
| <b>Step 1</b> | Click <b>Add Destination</b> :  |
| <b>Step 2</b> | Configure the entry.            |
| <b>Step 3</b> | Submit and commit your changes. |
- 

## Importing and Exporting Destination Control Configurations

If you are managing multiple domains, you can create a single configuration file to define Destination Control entries for all of the domains and import it onto the appliance. The format of the configuration file is similar to a Windows INI configuration file. The parameters for a domain are grouped in a section with the domain name as the section name. For example, use the section name `[example.com]` to group the parameters for the domain `example.com`. Any parameter that is not defined will be inherited from the default Destination Control entry. You can define the parameters for the default Destination Control entry by including a `[DEFAULT]` section in the configuration file.

Importing the configuration file overwrites all of appliance's Destination Control entries, except for the default entry unless the configuration file includes the `[DEFAULT]` section. All other existing Destination Control entries will be deleted.

You can define any of the following parameters for a domain in the configuration file. All parameters are required for the [DEFAULT] section except for the `bounce_profile` parameter:

**Table 24-9 Destination Control Configuration File Parameters**

| Parameter Name                           | Description                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>ip_sort_pref</code>                | Specifies the Internet Protocol version for the domain.<br>Enter one of the following values: <ul style="list-style-type: none"> <li><code>PREFER_V6</code> for “IPv6 Preferred”</li> <li><code>REQUIRE_V6</code> for “IPv6 Required”</li> <li><code>PREFER_V4</code> for “IPv4 Preferred”</li> <li><code>REQUIRE_V4</code> for “IPv4 Required”</li> </ul>   |
| <code>max_host_concurrency</code>        | The maximum number of outbound connections that will be made by the appliance to a given host.<br>If you define this parameter for a domain, the <code>limit_type</code> and <code>limit_apply</code> parameters must also be defined.                                                                                                                       |
| <code>max_messages_per_connection</code> | The maximum number of messages allowed for a single outbound connection from the appliance to a given host before initiating a new connection.                                                                                                                                                                                                               |
| <code>recipient_minutes</code>           | The period of time — between 1 and 60 minutes — that the appliance will count the number of recipients. Leave undefined if no recipient limit should be applied.                                                                                                                                                                                             |
| <code>recipient_limit</code>             | The maximum number of recipients allowed within the given period of time. Leave undefined if no recipient limit should be applied.<br>If you define this parameter for a domain, the <code>recipient_minutes</code> , <code>limit_type</code> , and <code>limit_apply</code> parameters must also be defined.                                                |
| <code>limit_type</code>                  | Specifies whether the limit will be applied to the entire domain or to each mail exchange IP address specified for that domain.<br>Enter one of the following values: <ul style="list-style-type: none"> <li><code>0</code> (or <code>host</code>) for the domain</li> <li><code>1</code> (or <code>MXIP</code>) for the mail exchange IP address</li> </ul> |
| <code>limit_apply</code>                 | Specifies whether the limit will be applied system-wide or for each Virtual Gateway address.<br>Enter one of the following values: <ul style="list-style-type: none"> <li><code>0</code> (or <code>system</code>) for system-wide</li> <li><code>1</code> (or <code>VG</code>) for Virtual Gateway</li> </ul>                                                |
| <code>bounce_validation</code>           | Specifies whether to turn on bounce validation address tagging.<br>Enter one of the following values: <ul style="list-style-type: none"> <li><code>0</code> (or <code>off</code>)</li> <li><code>1</code> (or <code>on</code>)</li> </ul>                                                                                                                    |

**Table 24-9 Destination Control Configuration File Parameters**

| Parameter Name     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| table_tls          | <p>Specifies the TLS setting for the domain. See <a href="#">Enabling TLS and Certificate Verification on Delivery</a>, page 23-10 for more information.</p> <p>Enter one of the following values:</p> <ul style="list-style-type: none"> <li>• 0 (or off)</li> <li>• 1 (or on) for “Preferred”</li> <li>• 2 (or required) for “Required”</li> <li>• 3 (or on_verify) for “Preferred (Verify)”</li> <li>• 4 (or require_verify) for “Required (Verify)”</li> </ul> <p>Strings are not case sensitive.</p> |
| bounce_profile     | Name of the bounce profile to use. This cannot be used in the [DEFAULT] destination control entry.                                                                                                                                                                                                                                                                                                                                                                                                        |
| send_tls_req_alert | <p>Whether to send an alert if the required TLS connection fails.</p> <p>Enter one of the following values:</p> <ul style="list-style-type: none"> <li>• 0 (or off)</li> <li>• 1 (or on)</li> </ul> <p>This is a global setting and can only be used in the [DEFAULT] destination control entry.</p>                                                                                                                                                                                                      |
| certificate        | <p>Certificate used for outgoing TLS connections. This is a global setting and can only be used in the [DEFAULT] destination control entry.</p> <p><b>Note</b> If you do not specify a certificate, AsyncOS assigns the demonstration certificate, but using the demonstration certificate is not secure and not recommended for general use.</p>                                                                                                                                                         |

The following example shows a configuration file for the domains example1.com and example2.com along with the default Destination Control entry:

```
[DEFAULT]

ip_sort_pref = PREFER_V6

max_host_concurrency = 500

max_messages_per_connection = 50

recipient_minutes = 60

recipient_limit = 300

limit_type = host
```

```

limit_apply = VG

table_tls = off

bounce_validation = 0

send_tls_req_alert = 0

certificate = example.com

[example1.com]

ip_sort_pref = PREFER_V6

recipient_minutes = 60

recipient_limit = 100

table_tls = require_verify

limit_apply = VG

bounce_profile = tls_failed

limit_type = host

[example2.com]

table_tls = on

bounce_profile = tls_failed

```

The above example results in the following Destination Control entries for example1.com and example2.com:

```

example1.com

 IP Address Preference: IPv6 Preferred

 Maximum messages per connection: 50

 Rate Limiting:

 500 concurrent connections

 100 recipients per 60 minutes

 Limits applied to entire domain, across all virtual gateways

```

```

 TLS: Required (Verify)

 Bounce Profile: tls_failed

example2.com

 IP Address Preference: IPv6 Preferred

 Maximum messages per connection: Default

 Rate Limiting: Default

 TLS: Preferred

 Bounce Profile: tls_failed

```

Use the **Import Table** button on the Destination Controls page or the `destconfig -> import` command to import a configuration file. You can also export your Destination Control entries to an INI file using the **Export Table** button on the Destination Controls page or the `destconfig -> export` command. AsyncOS includes the `[Default]` domain control entry in the exported INI file.

## Destination Controls and the CLI

You can use the `destconfig` command in the CLI to configure Destination Control entries. This command is discussed in the *Cisco AsyncOS CLI Reference Guide*.

# Bounce Verification

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”).

In turn, for each incoming email with an invalid Envelope Recipient, the receiving mail servers generate a new email — a bounce message — and send it along to the Envelope Sender at the innocent domain (the one whose Envelope Sender address was forged). As a result, this target domain receives a flood of “misdirected” bounces — potentially millions of messages. This type of distributed denial of service attack can bring down email infrastructure and render it impossible for the target to send or receive legitimate email.

To combat these misdirected bounce attacks, AsyncOS includes Bounce Verification. When enabled, Bounce Verification tags the Envelope Sender address for messages sent via your appliance. The Envelope Recipient for any bounce message received by the 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.

Note that you can use Bounce Verification to manage incoming bounce messages based on your outgoing mail. To control how your appliance generates outgoing bounces (based on incoming mail), see [Directing Bounced Email, page 24-35](#).

#### Related Topics

- [Overview: Tagging and Bounce Verification, page 24-52](#)
- [Accepting Legitimate Untagged Bounced Messages, page 24-53](#)
- [Preventing a Bounced Message Storm Using Bounce Verification, page 24-54](#)

## Overview: Tagging and Bounce Verification

When sending email with bounce verification enabled, your appliance will rewrite the Envelope Sender address in the message. For example, MAIL FROM: joe@example.com becomes MAIL FROM: prvs=joe=123ABCDEFGFG@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 appliance. The tag is generated using a key defined in the Bounce Verification settings (see [Bounce Verification Address Tagging Keys, page 24-53](#) 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 (see [Working with Destination Controls, page 24-44](#)).

If a message already contains a tagged address, AsyncOS does not add another tag (in the case of an appliance delivering a bounce message to an appliance inside the DMZ).

#### Related Topics

- [Handling Incoming Bounce Messages, page 24-52](#)
- [Bounce Verification Address Tagging Keys, page 24-53](#)

## Handling Incoming Bounce Messages

Bounces that include a valid tag are delivered. The tag is removed and the Envelope Recipient is restored. This occurs immediately after the Domain Map step in the email pipeline. You can define how your appliances handle untagged or invalidly tagged bounces — reject them or add a custom header. See [Configuring Bounce Verification Settings, page 24-55](#) for more information.

If the bounce verification tag is not present, or if the key used to generate the tag has changed, or if the message is more than seven days old, the message is treated as per the settings defined for Bounce Verification.

For example, the following mail log shows a bounced message rejected by the appliance:

```
Fri Jul 21 16:02:19 2006 Info: Start MID 26603 ICID 125192
```

```
Fri Jul 21 16:02:19 2006 Info: MID 26603 ICID 125192 From: <>
```

```
Fri Jul 21 16:02:40 2006 Info: MID 26603 ICID 125192 invalid bounce, rcpt address
<bob@example.com> rejected by bounce verification.
```

```
Fri Jul 21 16:03:51 2006 Info: Message aborted MID 26603 Receiving aborted by sender
```

```
Fri Jul 21 16:03:51 2006 Info: Message finished MID 26603 aborted
```

**Note**

When delivering non-bounce mail to your own internal mail server (Exchange, etc.), you should disable Bounce Verification tagging for that internal domain.

AsyncOS considers bounces as mail with a null Mail From address (<>). For non-bounce messages that might contain a tagged Envelope Recipient, AsyncOS applies a more lenient policy. In such cases, AsyncOS ignores the seven-day key expiration and tries to find a match with older keys as well.

## Bounce Verification Address Tagging Keys

The tagging key is a text string your appliance uses when generating the bounce verification tag. Ideally, you would use the same key across all of your appliances so that all mail leaving your domain is tagged consistently. That way, if one appliance tags the Envelope Sender on an outgoing message an incoming bounce will be verified and delivered even if the bounce is received by a different appliance.

There is a seven day grace period for tags. For example, you may choose to change your tagging key multiple times within a seven-day period. In such a case, your appliance will try to verify tagged messages using all previous keys that are less than seven days old.

## Accepting Legitimate Untagged Bounced Messages

AsyncOS also includes a HAT setting related to Bounce Verification for considering whether untagged bounces are valid. The default setting is “No,” which means that untagged bounces are considered invalid and the appliance either rejects the message or applies a customer header, depending on the action selected on the Mail Policies > Bounce Verification page. If you select “Yes,” the appliance considers untagged bounces to be valid and accepts them. This may be used in the following scenario:

Suppose you have a user that wants to send email to a mailing list. However, the mailing list accepts messages only from a fixed set of Envelope Senders. In such a case, tagged messages from your user will not be accepted (as the tag changes regularly).

### Procedure

- Step 1** Add the domain to which the user is trying to send mail to the Destination Controls table and disable tagging for that domain. At this point, the user can send mail without problems.
- Step 2** However, to properly support receiving bounces from that domain (since they will not be tagged) you can create a sender group for that domain and enable the Consider Untagged Bounces to be Valid parameter in an “Accept” mail flow policy.



**Figure 24-6** The Consider Untagged Bounces to be Valid HAT Parameter

| Security Features                                                                                         |                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Spam Detection:                                                                                           | <input checked="" type="radio"/> Use Default (On) <input type="radio"/> On <input type="radio"/> Off                                                                                    |
| Virus Protection:                                                                                         | <input checked="" type="radio"/> Use Default (On) <input type="radio"/> On <input type="radio"/> Off                                                                                    |
| Encryption and Authentication:                                                                            | TLS: <input checked="" type="radio"/> Use Default (Off) <input type="radio"/> Off <input type="radio"/> Preferred <input type="radio"/> Required                                        |
|                                                                                                           | SMTP Authentication: <input checked="" type="radio"/> Use Default (Off) <input type="radio"/> Off <input type="radio"/> Preferred <input type="radio"/> Required                        |
|                                                                                                           | If Both TLS and SMTP Authentication are enabled: <input type="checkbox"/> Require TLS To Offer SMTP Authentication                                                                      |
| Domain Key/DKIM Signing:                                                                                  | <input checked="" type="radio"/> Use Default (Off) <input type="radio"/> On <input type="radio"/> Off                                                                                   |
| DKIM Verification:                                                                                        | <input checked="" type="radio"/> Use Default (Off) <input type="radio"/> On <input type="radio"/> Off                                                                                   |
| SPF/SIDF Verification:                                                                                    | <input checked="" type="radio"/> Use Default (Off) <input type="radio"/> On <input type="radio"/> Off                                                                                   |
|                                                                                                           | Conformance Level: <input type="text" value="SIDF Compatible"/>                                                                                                                         |
|                                                                                                           | Downgrade PRA verification result if 'Resent-Sender:' or 'Resent-From:' were used: <input checked="" type="radio"/> Use Default (No) <input type="radio"/> No <input type="radio"/> Yes |
|                                                                                                           | HELO Test: <input checked="" type="radio"/> Use Default (On) <input type="radio"/> Off <input type="radio"/> On                                                                         |
| Bounce Verification:                                                                                      | Consider Untagged Bounces to be Valid: <input checked="" type="radio"/> Use Default (No) <input type="radio"/> Yes <input type="radio"/> No                                             |
| (Applies only if bounce verification address tagging is in use. See Mail Policies > Bounce Verification.) |                                                                                                                                                                                         |

## Preventing a Bounced Message Storm Using Bounce Verification

### Procedure

- Step 1** Enter a tagging key. For more information, see [Configuring Bounce Verification Address Tagging Keys](#), page 24-55.
- Step 2** Edit the bounce verification settings. For more information, see [Configuring Bounce Verification Settings](#), page 24-55.
- Step 3** Enable bounce verification via Destination Controls. For more information, see [Working with Destination Controls](#), page 24-44.

**Figure 24-7** IronPort Bounce Verification Page

| Bounce Verification Settings         |         |
|--------------------------------------|---------|
| Action when invalid bounce received: | Reject  |
| Smart exceptions to tagging:         | Enabled |
| <a href="#">Edit Settings</a>        |         |

| Bounce Verification Address Tagging Keys                                      |                                                                                                                                          |
|-------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">Clear All Keys</a>                                                |                                                                                                                                          |
| <a href="#">New Key...</a>                                                    |                                                                                                                                          |
| Address Tagging Keys                                                          | Status                                                                                                                                   |
| example.com's bounce key                                                      | Current<br>(see Mail Policies > Destination Controls to set or view destinations which have Bounce Verification Address Tagging enabled) |
| <a href="#">Purge Keys</a> <input type="text" value="Not used in one month"/> |                                                                                                                                          |

Key:

### Related Topics

- [Configuring Bounce Verification Address Tagging Keys](#), page 24-55
- [Configuring Bounce Verification Settings](#), page 24-55
- [Configuring Bounce Verification Using the CLI](#), page 24-55
- [Bounce Verification and Cluster Configuration](#), page 24-55

## Configuring Bounce Verification Address Tagging Keys

The Bounce Verification Address Tagging Keys listing shows your current key and any unpurged keys you have used in the past. To add a new key:

### Procedure

- 
- |               |                                                                         |
|---------------|-------------------------------------------------------------------------|
| <b>Step 1</b> | On the Mail Policies > Bounce Verification page, click <b>New Key</b> . |
| <b>Step 2</b> | Enter a text string and click <b>Submit</b> .                           |
| <b>Step 3</b> | <b>Commit</b> your changes.                                             |
- 

### Related Topics

- [Purging Keys, page 24-55](#)

## Purging Keys

You can purge your old address tagging keys by selecting a rule for purging from the pull-down menu and clicking **Purge**.

## Configuring Bounce Verification Settings

The bounce verification settings determine which action to take when an invalid bounce is received.

### Procedure

- 
- |               |                                                                                                                                                                                                                                                                                      |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Choose Mail Policies > Bounce Verification.                                                                                                                                                                                                                                          |
| <b>Step 2</b> | Click <b>Edit Settings</b> .                                                                                                                                                                                                                                                         |
| <b>Step 3</b> | Select whether to reject invalid bounces, or to add a custom header to the message. If you want to add a header, enter the header name and value.                                                                                                                                    |
| <b>Step 4</b> | Optionally, enable smart exceptions. This setting allows incoming mail messages, and bounce messages generated by internal mail servers, to be automatically exempted from bounce verification processing (even when a single listener is used for both incoming and outgoing mail). |
| <b>Step 5</b> | Submit and commit your changes.                                                                                                                                                                                                                                                      |
- 

## Configuring Bounce Verification Using the CLI

You can use the `bvconfig` and `destconfig` commands in the CLI to configure bounce verification. These commands are discussed in the *Cisco AsyncOS CLI Reference Guide*.

## Bounce Verification and Cluster Configuration

Bounce verification works in a cluster configuration as long as both appliances use the same "bounce key." When you use the same key, either systems should be able to accept a legitimate bounce back. The modified header tag/key is not specific to each appliance.

# Set Email Delivery Parameters

The `deliveryconfig` command sets parameters to be used when delivering email from the appliance.

The appliance accepts email using multiple mail protocols: SMTP and QMQP. However, all outgoing email is delivered using SMTP, which is why the `deliveryconfig` command does not require that the protocol be specified.

**Note**

Several of the features or commands described in this section will affect, or be affected by routing precedence. Please see the “Assigning Network and IP Addresses” appendix for more information.

**Related Topics**

- [Default Delivery IP Interface, page 24-56](#)
- [Possible Delivery Feature, page 24-56](#)
- [Default Maximum Concurrency, page 24-57](#)
- [deliveryconfig Example, page 24-57](#)

## Default Delivery IP Interface

By default, the system uses an IP interface or IP interface group for email delivery. Any currently configured IP interface or IP interface group can be set. If no specific interface is identified, AsyncOS will use the hostname associated with the default delivery interface in the `SMTP HELO` command when communicating with recipient hosts. To configure IP interfaces, use the `interfaceconfig` command.

These are the rules for using Auto selection of email delivery interfaces:

- If the remote email server is on the same subnet as one of the configured interfaces, then traffic will go out on the matching interface.
- When set to auto-select, static routes you have configured using `routeconfig` take effect.
- Otherwise, the interface that is on the same subnet as the default gateway will be used. If all of the IP addresses have an equivalent route to the destination, then the system uses the most efficient interface available.

## Possible Delivery Feature

**Caution**

If you enable this feature, message delivery will not be reliable and may lead to loss of messages. Also, your appliance will not be RFC 5321-compliant. For more information, see <http://tools.ietf.org/html/rfc5321#section-6.1>.

When the Possible Delivery feature is enabled, AsyncOS treats any message that times-out after the body of the message is delivered, but before recipient host acknowledges receipt of the message, as a “possible delivery.” This functionality prevents recipients from receiving multiple copies of a message if continuous errors at their recipient host prevent acknowledgment of receipt. AsyncOS logs this recipient as a possible delivery in the mail logs and counts the message as completed.

## Default Maximum Concurrency

You also specify the default maximum number of concurrent connections the appliance makes for outbound message delivery. (The system-wide default is 10,000 connections to separate domains.) The limit is monitored in conjunction with the per-listener maximum outbound message delivery concurrency (the default per listener is 600 connections for private listeners and 1000 connections for public listeners). Setting the value lower than the default prevents the gateway from dominating weaker networks. For example, certain firewalls do not support large numbers of connections, and this could induce Denial of Service (DoS) warnings in these environments.

## deliveryconfig Example

In the following example, the `deliveryconfig` command is used to set the default interface to “Auto” with “Possible Delivery” enabled. The system-wide maximum outbound message delivery is set to 9000 connections.

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

```
Choose the operation you want to perform:
```

```
- SETUP - Configure mail delivery.
```

```
[> setup
```

```
Choose the default interface to deliver mail.
```

1. Auto
2. PublicNet2 (192.168.3.1/24: mail4.example.com)
3. Management (192.168.42.42/24: mail3.example.com)
4. PrivateNet (192.168.1.1/24: mail3.example.com)
5. PublicNet (192.168.2.1/24: mail3.example.com)

```
[1]> 1
```


```
Enable "Possible Delivery" (recommended)? [Y]> y
```

```
Please enter the default system wide maximum outbound message delivery
concurrency
```

```
[10000]> 9000
```

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

Our Email Gateway configuration now looks like this:

**Figure 24-8** *Setting Destination and Delivery Parameters*

A `destconfig` entry for the host `small-isp.net` was used to limit 100 simultaneous connections, or 10 simultaneous connections using Virtual Gateway addresses.

The `deliveryconfig` command was used to use Auto-selection of interfaces for email delivery and the Possible Delivery feature was enabled. The system-wide maximum outbound message delivery was set to 9000 total concurrent connections.

## Configuring Mail Gateways for all Hosted Domains Using Virtual Gateway™ Technology

This section describes Cisco Virtual Gateway™ technology and its benefits, how to set up a Virtual Gateway address, and how to monitor and manage Virtual Gateway addresses.

The Cisco Virtual Gateway technology allows you to configure enterprise mail gateways for all domains you host — with distinct IP addresses, hostname and domains — and create separate corporate email policy enforcement and anti-spam strategies for those domains, while hosted within the same physical appliance. The number of Virtual Gateway addresses available on all Email Security appliance models is 255.

**Related Topics**

- [Overview, page 24-60](#)
- [Setting Up Virtual Gateway Addresses, page 24-60](#)
- [Monitoring the Virtual Gateway Addresses, page 24-68](#)
- [Managing Delivery Connections per Virtual Gateway Address, page 24-68](#)

## Overview

Cisco has developed a unique Virtual Gateway technology designed to help ensure that corporations can reliably communicate with their customers via email. Virtual Gateway technology enables users to separate the appliance into multiple Virtual Gateway addresses from which to send and receive email. Each Virtual Gateway address is given a distinct IP address, hostname and domain, and email queue.

Assigning a distinct IP address and hostname to each Virtual Gateway address ensures that email delivered through the gateway will be properly identified by the recipient host and prevents critical email from being blocked as spam. The appliance has the intelligence to give the correct hostname in the `SMTP HELO` command for each of the Virtual Gateway addresses. This ensures that if a receiving Internet Service Provider (ISP) performs a reverse DNS look-up, the appliance will match the IP address of the email sent through that Virtual Gateway address. This feature is extremely valuable, because many ISPs use a reverse DNS lookup to detect unsolicited email. If the IP address in the reverse DNS look-up does not match the IP address of the sending host, the ISP may assume the sender is illegitimate and will frequently discard the email. The Cisco Virtual Gateway technology ensures that reverse DNS look-ups will always match the sending IP address, preventing messages from being blocked accidentally.

Messages in each Virtual Gateway address are also assigned to a separate message queue. If a certain recipient host is blocking email from one Virtual Gateway address, messages intended for that host will remain in the queue and eventually timeout. But messages intended for the same domain in a different Virtual Gateway queue that is not being blocked will be delivered normally. While these queues are treated separately for delivery purposes, the system administration, logging and reporting capability still provide a holistic view into all Virtual Gateway queues as if they were one.

## Setting Up Virtual Gateway Addresses

Before setting up the Cisco Virtual Gateway addresses, you must allocate a set of IP addresses that will be used to send email from. (For more information, see the “Assigning Network and IP Addresses” appendix.) You should also ensure proper configuration of your DNS servers so that the IP address resolves to a valid hostname. Proper configuration of DNS servers ensures that if the recipient host performs a reverse DNS lookup, it will resolve to valid IP/hostname pairs.

**Related Topics**

- [Creating New IP Interfaces for Use with Virtual Gateways, page 24-61](#)
- [Mapping Messages to IP Interfaces for Delivery, page 24-63](#)
- [Importing an `altsrchost` File, page 24-64](#)

- [altsrchost Limits](#), page 24-64
- [Example Text File with Valid Mappings for the altsrchost Command](#), page 24-65
- [Adding an altsrchost Mapping through the CLI](#), page 24-65

## Creating New IP Interfaces for Use with Virtual Gateways

After the IP addresses and hostnames have been established, the first step in configuring the Virtual Gateway addresses is to create new IP interfaces with the IP/hostname pairs using the Network > IP Interfaces page in the GUI or the `interfaceconfig` command in the CLI.

Once the IP interfaces have been configured, you have the option to combine multiple IP interfaces into interface groups; these groups can then be assigned to specific Virtual Gateways addresses which the system cycles through in a “round robin” fashion when delivering email.


After creating the required IP interfaces, you have two options for setting up the Virtual Gateway addresses and defining which email campaign will be sent from each IP interface or interface group:

- You can use the `altsrchost` command to map email from specific sender IP addresses or Envelope Sender address information to a host IP interface (Virtual Gateway address) or interface group for delivery.
- Using message filters, you can set up specific filters to deliver flagged messages using a specific host IP interface (Virtual Gateway address) or interface group. See [Alter Source Host \(Virtual Gateway address\) Action](#), page 9-68. (This method is more flexible and powerful than the one above.)

For more information about creating IP interfaces, see the “Accessing the Appliance” appendix.

So far, we have been using an Email Gateway configuration with the following interfaces defined as shown in [Figure 24-9](#).

**Figure 24-9**      **Example Public and Private Interfaces**



In the following example, the IP Interfaces page confirms that these two interfaces (PrivateNet and PublicNet) have been configured, in addition to the Management interface.



**Figure 24-10 IP Interfaces Page**  
**IP Interfaces**

| Network Interfaces and IP Addresses |                  |                   |        |
|-------------------------------------|------------------|-------------------|--------|
| <a href="#">Add IP Interface...</a> |                  |                   |        |
| Name                                | IP Address       | Hostname          | Delete |
| Management                          | 192.168.42.42/24 | mail3.example.com |        |
| PrivateNet                          | 192.168.1.1/24   | mail3.example.com |        |
| PublicNet                           | 192.168.2.1/24   | mail3.example.com |        |

Next, the Add IP Interface page is used to create a new interface named PublicNet2 on the Data2 Ethernet interface. The IP address of 192.168.2.2 is used, and the hostname of mail4.example.com is specified. The services for FTP (port 21), Telnet (port 23), and SSH (port 22) are then enabled.

**Figure 24-11 Add IP Interface Page**  
**Add IP Interface**

**IP Interface Settings**

Name:

Ethernet Port:

IP Address:  \*

Netmask:  \*

Hostname:


Services:

| Service                                                                                                                                                                                  | Port                               |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| <input checked="" type="checkbox"/> FTP                                                                                                                                                  | <input type="text" value="21"/>    |
| <input checked="" type="checkbox"/> Telnet                                                                                                                                               | <input type="text" value="23"/>    |
| <input checked="" type="checkbox"/> SSH                                                                                                                                                  | <input type="text" value="22"/> *  |
| <b>Appliance Management</b>                                                                                                                                                              |                                    |
| <input type="checkbox"/> HTTP                                                                                                                                                            | <input type="text" value="80"/> *  |
| <input type="checkbox"/> HTTPS                                                                                                                                                           | <input type="text" value="443"/> * |
| <input type="checkbox"/> Redirect HTTP requests to HTTPS (HTTP and HTTPS Services will be turned on)                                                                                     |                                    |
| <b>IronPort Spam Quarantine</b>                                                                                                                                                          |                                    |
| <input type="checkbox"/> IronPort Spam Quarantine HTTP                                                                                                                                   | <input type="text" value="82"/>    |
| <input type="checkbox"/> IronPort Spam Quarantine HTTPS                                                                                                                                  | <input type="text" value="93"/>    |
| <input type="checkbox"/> Redirect HTTP requests to HTTPS (HTTP and HTTPS Services will be turned on)                                                                                     |                                    |
| <input type="checkbox"/> This is the default interface for IronPort Spam Quarantine<br>Quarantine login and notifications will originate on this interface.                              |                                    |
| URL Displayed in Notifications:                                                                                                                                                          |                                    |
| <input checked="" type="radio"/> Hostname<br><input type="radio"/> <input type="text" value="http://10.1.1.1:82/"/><br><small>(examples: http://spamQ.url/, http://10.1.1.1:82/)</small> |                                    |


Warnings - \* Please exercise care when disabling or changing these items, as this could disrupt active connections to this appliance when changes to these items are committed.  
\*\* Hyperlinks and URLs affected by these changes will not be usable until the changes are committed.

Our Email Gateway configuration now looks like this:

**Figure 24-12 Adding Another Public Interface**



Using Virtual Gateway addresses, a configuration like the one shown in [Figure 24-13](#) is also possible.

**Figure 24-13** Four Virtual Gateway Addresses on One Ethernet Interface

Note that four separate IP interfaces can be used to deliver mail, where only one public listener is configured to accept messages from the Internet.

## Mapping Messages to IP Interfaces for Delivery

The `altsrchost` command provides the simplest and most straightforward method to segment each appliance into multiple IP interfaces (Virtual Gateway addresses) from which to deliver email. However, users requiring more power and flexibility in mapping messages to particular Virtual Gateways should investigate the use of message filters. See [Chapter 9, “Using Message Filters to Enforce Email Policies”](#) for more information.

The `altsrchost` command allows you to control which IP interface or interface group to use during email delivery based on one of the following:

- the sender’s IP address
- the Envelope Sender address

To specify which IP interface or interface group the system will deliver email from, you create mapping keys that pair either the sender’s IP address or the Envelope Sender address to an IP interface or interface group (specified by interface name or group name).

AsyncOS will compare both the IP address and Envelope Sender address to the mapping keys. If either the IP address or Envelope Sender address matches one of the keys, the corresponding IP interface is used for the outbound delivery. If there is no match, the default outbound interface will be used.

The system can match any of the following keys and take preference in the following order:

|                                     |                                                                                             |
|-------------------------------------|---------------------------------------------------------------------------------------------|
| <b>Sender’s IP address</b>          | The IP address of the sender must match exactly.<br>Example: 192.168.1.5                    |
| <b>Fully-formed Envelope Sender</b> | The Envelope Sender must match the entire address exactly.<br>Example: username@example.com |

|                 |                                                                                                                                                                        |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Username</b> | The system will match username syntax against the Envelope Sender address up to the @ sign. The @ sign must be included. Example: <code>username@</code>               |
| <b>Domain</b>   | The system will match domain name syntax against the Envelope Sender address starting with the @ sign. The @ sign must be included. Example: <code>@example.com</code> |

**Note**

A listener checks the information in the `altsrchost` table and directs the email to a particular interface *after* checking the masquerading information and *before* message filters are checked.

Use these subcommands within the `altsrchost` command to create mappings in the Virtual Gateways via the CLI:

| Syntax        | Description                                |
|---------------|--------------------------------------------|
| <b>new</b>    | Create a new mapping manually.             |
| <b>print</b>  | Display the current list of mappings.      |
| <b>delete</b> | Remove one of the mappings from the table. |

## Importing an altsrchost File

Like the HAT, the RAT, `smtproutes`, and masquerading and alias tables, you can modify `altsrchost` entries by exporting and importing a file.

### Procedure

- Step 1** Use the `export` subcommand of the `altsrchost` command to export the existing entries to a file (whose name you specify).
- Step 2** Outside of the CLI, get the file. (See [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) for more information.)
- Step 3** With a text editor, create new entries in the file. The order that rules appear in the `altsrchost` table is important.
- Step 4** Save the file and place it in the “altsrchost” directory for the interface so that it can be imported. (See [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) for more information.)
- Step 5** Use the `import` subcommand of `altsrchost` to import the edited file.

## altsrchost Limits

You can define up to 1,000 `altsrchost` entries.

## Example Text File with Valid Mappings for the altsrchoost Command

```
Comments to describe the file

@example.com DemoInterface

paul@ PublicInterface

joe@ PublicInterface

192.168.1.5, DemoInterface

steve@example.com PublicNet
```

The `import` and `export` subcommands operate on a line-by-line basis and map either the sender IP address or the Envelope Sender address line to the interface name. The key must be the first block of non-space characters followed by the interface name in the second block of non-space characters, separated by a comma (,) or space (.). Comment lines start with a number sign (#) and will be ignored.

## Adding an altsrchoost Mapping through the CLI

In the following example, the `altsrchoost` table is printed to show that there are no existing mappings. Two entries are then created:

- Mail from the groupware server host named `@exchange.example.com` is mapped to the `PublicNet` interface.
- Mail from the sender IP address of `192.168.35.35` (for example, the marketing campaign messaging system) is mapped to the `PublicNet` interface.

Finally, the `altsrchoost` mappings are printed to confirm and the changes are committed.

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

```
There are currently no mappings configured.
```

```
Choose the operation you want to perform:
```

- NEW - Create a new mapping.
- IMPORT - Load new mappings from a file.

```
[> new
```

```
Enter the Envelope From address or client IP address for which you want to set up a
Virtual Gateway mapping. Partial addresses such as "@example.com" or "user@" are
allowed.
```

```
[> @exchange.example.com
```

Which interface do you want to send messages for @exchange.example.com from?

1. PublicNet2 (192.168.2.2/24: mail4.example.com)
2. Management (192.168.42.42/24: mail3.example.com)
3. PrivateNet (192.168.1.1/24: mail3.example.com)
4. PublicNet (192.168.2.1/24: mail4.example.com)

[1]> **4**

Mapping for @exchange.example.com on interface PublicNet created.

Choose the operation you want to perform:

- NEW - Create a new mapping.
- EDIT - Modify a mapping.
- DELETE - Remove a mapping.
- IMPORT - Load new mappings from a file.
- EXPORT - Export all mappings to a file.
- PRINT - Display all mappings.
- CLEAR - Remove all mappings.

[> **new**

Enter the Envelope From address or client IP address for which you want to set up a Virtual Gateway mapping. Partial addresses such as "@example.com" or "user@" are allowed.

[> **192.168.35.35**

Which interface do you want to send messages for 192.168.35.35 from?

1. PublicNet2 (192.168.2.2/24: mail4.example.com)
2. Management (192.168.42.42/24: mail3.example.com)
3. PrivateNet (192.168.1.1/24: mail3.example.com)
4. PublicNet (192.168.2.1/24: mail4.example.com)

```
[1]> 1
```

Mapping for 192.168.35.35 on interface PublicNet2 created.

Choose the operation you want to perform:

- NEW - Create a new mapping.
- EDIT - Modify a mapping.
- DELETE - Remove a mapping.
- IMPORT - Load new mappings from a file.
- EXPORT - Export all mappings to a file.
- PRINT - Display all mappings.
- CLEAR - Remove all mappings.

```
[> print
```

1. 192.168.35.35 -> PublicNet2
2. @exchange.example.com -> PublicNet

Choose the operation you want to perform:

- NEW - Create a new mapping.
- EDIT - Modify a mapping.
- DELETE - Remove a mapping.
- IMPORT - Load new mappings from a file.
- EXPORT - Export all mappings to a file.
- PRINT - Display all mappings.
- CLEAR - Remove all mappings.

```
[>
```

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

Please enter some comments describing your changes:


```
[> Added 2 altsrchost mappings
```

Do you want to save the current configuration for rollback? [Y]> n

Changes committed: Fri May 23 11:42:12 2014 GMT

An illustration of the configuration change in this example is shown in Figure 24-14:

**Figure 24-14 Example: Selecting an IP Interface or Interface Group to Use**



## Monitoring the Virtual Gateway Addresses

While each Virtual Gateway address has its own email queue for delivery purposes, the system administration, logging, and reporting capabilities still provide a holistic view into all Virtual Gateway queues as if they were one. To monitor the recipient host status for each Virtual Gateway queue, use the `hoststatus` and `hostrate` command. See the “Reading the Available Components of Monitoring” section in the “Managing and Monitoring Using the CLI” chapter.

The `hoststatus` command returns monitoring information about email operations relating to a specific recipient host.

If you are using Virtual Gateway technology, information about each Virtual Gateway address is also displayed. The command requires you to input the domain of the host information to be returned. DNS information stored in the AsyncOS cache and the last error returned from the recipient host is also given. Data returned is cumulative since the last `resetcounters` command.

The statistics returned are grouped into two categories: counters and gauges. In addition, other data returned include: last activity, MX records, and last 5XX error.

## Managing Delivery Connections per Virtual Gateway Address

Certain system parameters require settings at the system and Virtual Gateway address levels.

For example, some recipient ISPs limit the number of connections they allow for each client host. Therefore, it is important to manage relationships with the ISPs, especially when email is being delivered over multiple Virtual Gateway addresses.

See [Controlling Email Delivery Using Destination Controls, page 24-42](#) for information about the `destconfig` command and how Virtual Gateway addresses are affected.

When you create a “group,” of Virtual Gateway addresses, the good neighbor table settings for Virtual Gateway are applied to the group, even if the group consists of 254 IP addresses.

For example, suppose you have created group of 254 outbound IP addresses set up as a group to cycle through in a “round-robin” fashion, and suppose the good neighbor table for `small-isp.com` is 100 simultaneous connections for the system and 10 connections for Virtual Gateway addresses. This configuration will *never* open more than 10 connections total for all 254 IP addresses in that group; the group is treated as a single Virtual Gateway address.

## Using Global Unsubscribe

To ensure that specific recipients, recipient domains, or IP addresses never receive messages from the appliance, use the AsyncOS Global Unsubscribe feature. The `unsubscribe` command allows you to add and delete addresses to a global unsubscribe list, as well as enable and disable the feature. AsyncOS checks all recipient addresses against a list of “globally unsubscribed” users, domains, email addresses, and IP addresses. If a recipient matches an address in the list, the recipient is either dropped or hard bounced, and the Global Unsubscribe (GUS) counter is incremented. (Log files will note whether a matching recipient was dropped or hard bounced.) The GUS check occurs immediately before an attempt to send email to a recipient, thus inspecting all messages sent by the system.



### Note

Global Unsubscribe is not intended to replace the removal of names and general maintenance of mailing lists. The feature is intended to act as a fail-safe mechanism to ensure email does not get delivered to inappropriate entities.

The global unsubscribe feature applies to private and public listeners.

Global Unsubscribe has a maximum limit of 10,000 addresses. To increase this limit, contact your Cisco sales representative. Global Unsubscribe addresses can be in one of four forms:

**Table 24-10 Global Unsubscribe Syntax**

|                                   |                                                                                                                                                              |
|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>username@example.com</code> | Fully-formed email address<br>This syntax is used to block a specific recipient at a specific domain.                                                        |
| <code>username@</code>            | Username<br>The username syntax will block all recipients with the specified username at all domains. The syntax is the username followed by an at sign (@). |
| <code>@example.com</code>         | Domain<br>The domain syntax is used to block all recipients destined for a particular domain. The syntax is the specific domain, preceded by an at sign (@). |



**Table 24-10**     *Global Unsubscribe Syntax (continued)*

|                            |                                                                                                                                                                                                                                                      |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>@.example.com</code> | Partial Domain<br><br>The partial domain syntax is used to block all recipients destined for a particular domain and all its subdomains.                                                                                                             |
| <code>10.1.28.12</code>    | IP address<br><br>The IP address syntax is used to block all recipients destined for a particular IP address. This syntax can be useful if a single IP address is hosting multiple domains. The syntax consists of a common dotted octet IP address. |

**Related Topics**

- [Adding a Global Unsubscribe Address Using The CLI, page 24-70](#)
- [Exporting and Importing a Global Unsubscribe File, page 24-72](#)

## Adding a Global Unsubscribe Address Using The CLI

In this example, the address `user@example.net` is added to the Global Unsubscribe list, and the feature is configured to hard bounce messages. Messages sent to this address will be bounced; the appliance will bounce the message immediately prior to delivery.

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

```
Global Unsubscribe is enabled. Action: drop.
```

```
Choose the operation you want to perform:
```

- ```
- NEW - Create a new entry.
- IMPORT - Import entries from a file.
- SETUP - Configure general settings.
```

```
[> new
```

```
Enter the unsubscribe key to add. Partial addresses such as
```

```
"@example.com" or "user@" are allowed, as are IP addresses. Partial hostnames such as
"@.example.com" are allowed.
```

```
[> user@example.net
```

```
Email Address 'user@example.net' added.
```

Global Unsubscribe is enabled.

Choose the operation you want to perform:

- NEW - Create a new entry.
- DELETE - Remove an entry.
- PRINT - Display all entries.
- IMPORT - Import entries from a file.
- EXPORT - Export all entries to a file.
- SETUP - Configure general settings.
- CLEAR - Remove all entries.

[> **setup**

Do you want to enable the Global Unsubscribe feature? [Y]> **y**

Would you like matching messages to be dropped or bounced?

1. Drop
2. Bounce

[1]> **2**

Global Unsubscribe is enabled. Action: bounce.

Choose the operation you want to perform:

- NEW - Create a new entry.
- DELETE - Remove an entry.
- PRINT - Display all entries.
- IMPORT - Import entries from a file.
- EXPORT - Export all entries to a file.
- SETUP - Configure general settings.
- CLEAR - Remove all entries.

```
[ ]>
```

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

```
Please enter some comments describing your changes:
```

```
[ ]> Added username "user@example.net" to global unsubscribe
```

```
Do you want to save the current configuration for rollback? [Y]> n
```

```
Changes committed: Fri May 23 11:42:12 2014 GMT
```

Exporting and Importing a Global Unsubscribe File

Like the HAT, the RAT, `smtproutes`, static masquerading tables, alias tables, domain map tables, and `altsrchost` entries, you can modify global unsubscribe entries by exporting and importing a file.

Procedure

-
- Step 1** Use the `export` subcommand of the `unsubscribe` command to export the existing entries to a file (whose name you specify).
 - Step 2** Outside of the CLI, get the file. (See [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) for more information.)
 - Step 3** With a text editor, create new entries in the file.
 Separate entries in the file by new lines. Return representations from all standard operating systems are acceptable (`<CR>`, `<LF>`, or `<CR><LF>`). Comment lines start with a number sign (`#`) and are ignored. For example, the following file excludes a single recipient email address (`test@example.com`), all recipients at a particular domain (`@testdomain.com`), all users with the same name at multiple domains (`testuser@`), and any recipients at a specific IP address (`11.12.13.14`).

```
# this is an example of the global_unsubscribe.txt file

test@example.com

@testdomain.com


testuser@

11.12.13.14
```
 - Step 4** Save the file and place it in the configuration directory for the interface so that it can be imported. (See [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) for more information.)

Step 5 Use the `import` subcommand of `unsubscribe` to import the edited file.

Our Email Gateway configuration now looks like this:

Figure 24-15 Global Unsubscribe Example



Review: Email Pipeline

Table 24-11 and Table 24-12 provide an overview of how email is routed through the system, from reception to routing to deliver. Each feature is processed in order (from top to bottom) and is briefly summarized. Shaded areas in Figure 24-15 represent processing that occurs in the Work Queue.

You can test most of the configurations of features in this pipeline using the `trace` command. For more information, see “Debugging Mail Flow Using Test Messages: Trace” in the Troubleshooting chapter.

**Note**

For outgoing mail, RSA Email Data Loss Prevention scanning takes place after the Outbreak Filters stage.

Table 24-11 *Email Pipeline for the Email Security 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.
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. <code>aliasconfig</code> is not a subcommand of <code>listenerconfig</code> .)
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.

Table 24-12 Email Pipeline for the Email Security 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.
	Anti-Spam**	Per Recipient Scanning	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.
	Advanced Malware Protection		Advanced Malware Protection performs file reputation scanning and file analysis, in order to detect malware in attachments.
	Content Filters*		Content Filters are applied. * Can send messages to quarantines.
	Outbreak Filters*		The Outbreak Filters feature helps protect against virus outbreaks. * Can send messages to quarantines.
	Virtual gateways		Sends mail over particular IP interfaces or groups of IP interfaces.
	Delivery limits		1. Sets the default delivery interface. 2. Sets the total maximum number of outbound connections.
	Domain-based Limits		Defines, per-domain: maximum outbound connections for each virtual gateway and for the entire system; the bounce profile to use; the TLS preference for delivery: no/preferred/required
	Domain-based routing		Routes mail based on domain without rewriting Envelope Recipient.
	Global unsubscribe		Drops recipients according to specific list (configured system-wide).
Bounce profiles		Undeliverable message handling. Configurable per listener, per Destination Controls entry, and via message filters.	

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



LDAP Queries

- [Overview of LDAP Queries, page 25-1](#)
- [Working with LDAP Queries, page 25-12](#)
- [Using Acceptance Queries For Recipient Validation, page 25-19](#)
- [Using Routing Queries to Send Mail to Multiple Target Addresses, page 25-20](#)
- [Using Masquerading Queries to Rewrite the Envelope Sender, page 25-21](#)
- [Using Group LDAP Queries to Determine if a Recipient is a Group Member, page 25-23](#)
- [Using Domain-based Queries to Route to a Particular Domain, page 25-26](#)
- [Using Chain Queries to Perform a Series of LDAP Queries, page 25-28](#)
- [Using LDAP For Directory Harvest Attack Prevention, page 25-29](#)
- [Configuring AsyncOS for SMTP Authentication, page 25-32](#)
- [Configuring External LDAP Authentication for Users, page 25-40](#)
- [Authenticating End-Users of the Spam Quarantine, page 25-43](#)
- [Spam Quarantine Alias Consolidation Queries, page 25-44](#)
- [Identifying a Sender's User Distinguished Name for RSA Enterprise Manager, page 25-45](#)
- [Configuring AsyncOS To Work With Multiple LDAP Servers, page 25-46](#)

Overview of LDAP Queries

If you store user information within LDAP directories in your network infrastructure — for example, in Microsoft Active Directory, SunONE Directory Server, or OpenLDAP directories — you can configure the appliance to query your LDAP servers to accept, route, and authenticate messages. You can configure the appliance to work with one or multiple LDAP servers.

The following section provides an overview on the types of LDAP queries you can perform; how LDAP works with the appliance to authenticate, accept, and route messages; and how to configure your appliance to work with LDAP.

Related Topics

- [Understanding LDAP Queries, page 25-2](#)
- [Understanding How LDAP Works with AsyncOS, page 25-3](#)
- [Configuring the Cisco IronPort Appliance to Work with an LDAP Server, page 25-4](#)

- [Creating LDAP Server Profiles to Store Information About the LDAP Server, page 25-5](#)
- [Testing LDAP Servers, page 25-6](#)
- [Enabling LDAP Queries to Run on a Particular Listener, page 25-7](#)
- [Enhanced Support for Microsoft Exchange 5.5, page 25-9](#)

Understanding LDAP Queries

If you store user information within LDAP directories in your network infrastructure, you can configure the appliance to query your LDAP server for the following purposes:

- **Acceptance Queries.** You can use your existing LDAP infrastructure to define how the recipient email address of incoming messages (on a public listener) should be handled. For more information, see [Using Acceptance Queries For Recipient Validation, page 25-19](#).
- **Routing (Aliasing).** You can configure the appliance to route messages to the appropriate address and/or mail host based upon the information available in LDAP directories on your network. For more information, see [Using Routing Queries to Send Mail to Multiple Target Addresses, page 25-20](#).
- **Certificate Authentication.** You can create a query that checks the validity of a client certificate in order to authenticate an SMTP session between the user's mail client and the Email Security appliance. For more information, see [Checking the Validity of a Client Certificate, page 26-51](#).
- **Masquerading.** You can masquerade Envelope Senders (for outgoing mail) and message headers (for incoming mail, such as To:, Reply To:, From: or CC:). For more information about masquerading, see [Using Masquerading Queries to Rewrite the Envelope Sender, page 25-21](#).
- **Group Queries.** You can configure the appliance to perform actions on messages based on the groups in the LDAP directory. You do this by associating a group query with a message filter. You can perform any message action available for message filters on messages that match the defined LDAP group. For more information, see [Using Group LDAP Queries to Determine if a Recipient is a Group Member, page 25-23](#).
- **Domain-based Queries.** You can create domain-based queries to allow the appliance to perform different queries for different domains on a single listener. When the Email Security Appliance runs the domain-based queries, it determines the query to use based on the domain, and it queries the LDAP server associated with that domain.
- **Chain Queries.** You can create a chain query to enable the appliance to perform a series of queries in sequence. When you configure a chain query, the appliance runs each query in sequence until the LDAP appliance returns a positive result.
- **Directory Harvest Prevention.** You can configure the appliance to combat directory harvest attacks using your LDAP directories. You can configure directory harvest prevention during the SMTP conversation or within the work queue. If the recipient is not found in the LDAP directory, you can configure the system to perform a delayed bounce or drop the message entirely. Consequently, spammers are not able to differentiate between valid and invalid email addresses. See [Using LDAP For Directory Harvest Attack Prevention, page 25-29](#).
- **SMTP Authentication.** AsyncOS provides support for SMTP authentication. SMTP Auth is a mechanism for authenticating clients connected to an SMTP server. You can use this functionality to enable users at your organization to send mail using your mail servers even if they are connecting remotely (e.g. from home or while traveling). For more information, see [Configuring AsyncOS for SMTP Authentication, page 25-32](#).


- **External Authentication.** You can configure your appliance to use your LDAP directory to authenticate users logging in to the appliance. For more information, see [Configuring External LDAP Authentication for Users](#), page 25-40.
- **Spam Quarantine End-User Authentication.** You can configure your appliance to validate users when they log in to the end-user quarantine. For more information, see [Authenticating End-Users of the Spam Quarantine](#), page 25-43.
- **Spam Quarantine Alias Consolidation.** If you use email notifications for spam, this query consolidates the end-user aliases so that end-users do not receive quarantine notices for each aliased email address. For more information, see [Spam Quarantine Alias Consolidation Queries](#), page 25-44.
- **User Distinguished Name.** If you use RSA Enterprise Manager for data loss prevention (DLP), this query retrieves the distinguished name for senders of messages that may contain DLP violations. The Email Security appliance includes the distinguished name when it sends DLP incident data to Enterprise Manager. For more information, see [Identifying a Sender's User Distinguished Name for RSA Enterprise Manager](#), page 25-45.

Understanding How LDAP Works with AsyncOS

When you work with LDAP directories, the appliance can be used in conjunction with an LDAP directory server to accept recipients, route messages, and/or masquerade headers. LDAP group queries can also be used in conjunction with message filters to create rules for handling messages as they are received by the appliance.

Figure 25-1 demonstrates how the appliance works with LDAP:

Figure 25-1 LDAP Configuration



1. The sending MTA sends a message to the public listener "A" via SMTP.
2. The appliance queries the LDAP server defined via the System Administration > LDAP page (or by the global `ldapconfig` command).

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
 - further actions as defined by `rcpt-to-group` or `mail-from-group` message filter rules (used in conjunction with configured group queries).

**Note**

You can configure your appliance to connect to multiple LDAP servers. When you do this, you can configure the LDAP profile settings for load-balancing or failover. For more information about working with multiple LDAP servers, see [Configuring AsyncOS To Work With Multiple LDAP Servers, page 25-46](#).

Configuring the Cisco IronPort Appliance to Work with an LDAP Server

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

Procedure

- Step 1 Configure LDAP server profiles.** The server profile contains information to enable AsyncOS 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

For more information about configuring a server profile, see [Creating LDAP Server Profiles to Store Information About the LDAP Server, page 25-5](#).

When you configure the LDAP server profile, you can configure AsyncOS to connect to one or multiple LDAP servers.

For information about configuring AsyncOS to connect to multiple servers, see [Configuring AsyncOS To Work With Multiple LDAP Servers, page 25-46](#).
- Step 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 [Understanding LDAP Queries, page 25-2](#).

For information on writing queries, see [Working with LDAP Queries, page 25-12](#).
- Step 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 [Enabling LDAP Queries to Run on a Particular Listener, page 25-7](#).

**Note**

When you configure a group query, you need to take additional steps to configure AsyncOS to work with the LDAP server. For information on configuring a group query, see [Using Group LDAP Queries to Determine if a Recipient is a Group Member, page 25-23](#). 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 the Spam Quarantine chapter.

Creating LDAP Server Profiles to Store Information About the LDAP Server

When you configure AsyncOS to use LDAP directories, you create an LDAP server profile to store the information about the LDAP server.

Procedure

- Step 1** On the System Administration > LDAP page, click **Add LDAP Server Profile**.
 - Step 2** Enter a name for the server profile.
 - Step 3** Enter the host name for the LDAP server.

You can enter multiple host names to configure the LDAP servers for failover or load-balancing. Separate multiple entries with commas. For more information, see [Configuring AsyncOS To Work With Multiple LDAP Servers, page 25-46](#).
 - Step 4** Select an authentication method. You can use anonymous authentication or specify a username and password.
 - Step 5** Select the LDAP server type: Active Directory, OpenLDAP, or Unknown or Other.
 - Step 6** Enter a port number.

The default port is 3268. This is the default port for Active Directory that enables it to access the global catalog in a multi-server environment.
 - Step 7** Enter a Base DN (distinguishing name) for the LDAP server.

If you authenticate with a username and a password, the username must include the full DN to the entry that contains the password. For example, a user is a member of the marketing group with an email address of joe@example.com. The entry for this user would look like the following entry:

```
uid=joe, ou=marketing, dc=example dc=com
```
 - Step 8** Select whether to use SSL when communicating with the LDAP server.
 - Step 9** Under Advanced, enter cache time-to-live. This value represents the amount of time to retain caches.
 - Step 10** Enter the maximum number of retained cache entries.
- **Note** This cache is maintained per LDAP server. If you are configuring more than one LDAP servers, you must set a smaller LDAP cache value for better performance. Also, if the memory usage of various processes in the appliance is high, increasing this value may reduce the system performance.
- Step 11** Enter a maximum number of simultaneous connections.

If you configure the LDAP server profile for load balancing, these connections are distributed among the listed LDAP servers. For example, if you configure 10 simultaneous connections and load balance the connections over three servers, AsyncOS creates 10 connections to each server, for a total of 30 connections.



Note The maximum number of simultaneous connections includes LDAP connections used for LDAP queries. However, the appliance may open more connections if you use LDAP authentication for the Spam Quarantine.

Step 12 Test the connection to the server by clicking the **Test Server(s)** button. If you specified multiple LDAP servers, they are all tested. The results of the test appear in the Connection Status field. For more information, see [Testing LDAP Servers, page 25-6](#).

Step 13 Create queries by marking the checkbox and completing the fields. You can select Accept, Routing, Masquerade, Group, SMTP Authentication, External Authentication, Spam Quarantine End-User Authentication, and Spam Quarantine Alias Consolidation.



Note To allow the appliance to run LDAP queries when you receive or send messages, you must enable the LDAP query on the appropriate listener. For more information, see [Enabling LDAP Queries to Run on a Particular Listener, page 25-7](#).

Step 14 Test a query by clicking the **Test Query** button.
Enter the test parameters and click **Run Test**. The results of the test appear in the Connection Status field. If you make any changes to the query definition or attributes, click **Update**. For more information, see [Testing LDAP Queries, page 25-17](#).



Note If you have configured the LDAP server to allow binds with empty passwords, the query can pass the test with an empty password field.

Step 15 Submit and commit your changes.



Note Although the number of server configurations is unlimited, you can configure only one recipient acceptance, one routing, one masquerading, and one group query per server.

Testing LDAP Servers

Use the **Test Server(s)** button on the Add/Edit LDAP Server Profile page (or the `test` subcommand of the `ldapconfig` command in the CLI) to test the connection to the LDAP server. AsyncOS displays a message stating whether the connection to the server port succeeded or failed. If you configured multiple LDAP servers, AsyncOS tests each server and displays individual results.

Enabling LDAP Queries to Run on a Particular Listener

To allow the appliance to run LDAP queries when you receive or send messages, you must enable the LDAP query on the appropriate listener.

Related Topics

- [Configuring Global Settings for LDAP Queries, page 25-7](#)
- [Example of Creating an LDAP Server Profile, page 25-7](#)
- [Enabling LDAP Queries on a Public Listener, page 25-8](#)
- [Enabling LDAP Queries on a Private Listener, page 25-9](#)

Configuring Global Settings for LDAP Queries

The LDAP global settings define how the appliance handles all LDAP traffic.

Procedure

-
Step 1	On the System Administration > LDAP page, click Edit Settings .
Step 2	Select the IP interface to use for LDAP traffic. The appliance automatically chooses an interface by default.
Step 3	Select the TLS certificate to use for the LDAP interface (TLS certificates added via the Network > Certificates page or the <code>certconfig</code> command in the CLI are available in the list, see Overview of Encrypting Communication with Other MTAs, page 23-1).
Step 4	Submit and commit your changes.
-

Example of Creating an LDAP Server Profile

In the following example, the System Administration > LDAP page is used to define an LDAP server for the appliance to bind to, and queries for recipient acceptance, routing, and masquerading are configured.



Note

There is a 60 second connection attempt time-out for LDAP connections (which covers the DNS lookup, the connection itself, and, if applicable, the authentication bind for the appliance itself). After the first failure, AsyncOS immediately starts trying other hosts in the same server (if you specified more than one in the comma separated list). If you only have one host in the server, AsyncOS continues attempting to connect to it.

Figure 25-2 Configuring an LDAP Server Profile (1 of 2)

LDAP Server Settings	
Server Attributes	
LDAP Server Profile Name:	PublicLDAP
Host Name(s):	myldapserver.example.com Fully qualified hostname or IP, separate multiple entries with a comma
Authentication Method:	☐ Anonymous ☒ Use Password Username: cn=anonymous Password: *****
Server Type: ?	Active Directory
Port: ?	3268
Base DN: ?	dc=example, dc=com
Connection Protocol:	☐ Use SSL
Advanced:	Cache TTL (time-to-live): 900 Seconds Maximum Retained Cache Entries: 10000 Maximum number of simultaneous connections for each host: 10 Multiple host options: ☒ Load-balance connections among all hosts listed ☐ Failover connections in the order listed
Server Attribute Testing:	Test Server(s)

First, the nickname of “PublicLDAP” is given for the myldapserver.example.com LDAP server. The number of connections is set to 10 (the default), and the multiple LDAP server (hosts) load balance option is left as the default. You can specify multiple hosts here by providing a comma separated list of names. Queries are directed to port 3268 (the default). SSL is not enabled as the connection protocol for this host. The base DN of example.com is defined (dc=example, dc=com). The cache time-to-live is set to 900 seconds, the maximum number of cache entries is 10000, and the authentication method is set to password.

Queries for recipient acceptance, mail routing, and masquerading are defined. Remember that query names are case-sensitive and must match exactly in order to return the proper results.

Figure 25-3 Configuring an LDAP Server Profile (2 of 2)

☒ Accept Query	
Name:	PublicLDAP.accept
Query String:	{proxyAddresses=smtp:{a}} Test Query
☒ Routing Query	
Name:	PublicLDAP.routing
Query String:	{mailLocalAddress={a}} Test Query
Recipient Email to Rewrite the Envelope Header:	mailRoutingAddress
Alternative Mailhost Attribute:	mailHost
SMTP Call-Ahead Server Attribute (optional):	This attribute is used only if an SMTP Call-Ahead server is configured. Go to Network > SMTP Call-Ahead.
☒ Masquerade Query	
Name:	PublicLDAP.masquerade
Query String:	{mailRoutingAddress={a}} Test Query
Attribute Containing Externally Visible Full Email Address:	mailLocalAddress

Enabling LDAP Queries on a Public Listener

In this example, the public listener “InboundMail” is updated to use LDAP queries for recipient acceptance. Further, recipient acceptance is configured to happen during the SMTP conversation (for more information, see [Using Acceptance Queries For Recipient Validation](#), page 25-19 for more information).

Figure 25-4 Enabling Acceptance and Routing Queries on a Listener

LDAP Queries: ▾ Accept

Accept Query: exampleTest.accept ▾

☐ Work Queue

Non-Matching Recipients: Bounce ▾

☒ SMTP Conversation

If the LDAP server is unreachable:

☐ Allow Mail in

☒ Drop Connection, return error code:

Code: 451

Text: Temporary recipient validation er

When the Directory Harvest Attack Prevention threshold (maximum invalid recipients per hour) is reached:

Code: 550

Text: Too many invalid recipients

☒ Drop Connection if the Directory Harvest Attack Prevention threshold (maximum invalid recipients per hour) is reached within an SMTP conversation.

▸ Routing

▸ Masquerade

▸ Group

Enabling LDAP Queries on a Private Listener

In this example, the private listener “OutboundMail” is updated to use LDAP queries for masquerading. The masqueraded fields include: From, To, CC, and Reply-To.

Figure 25-5 Enabling a Masquerading Query on a Listener

LDAP Queries: ▾ Accept

▸ Routing

▾ Masquerade

Masquerade Query: exampleTest.masquerade ▾

Addresses to Masquerade:

☐ Envelope Sender

☒ From (Header)

☒ To (Header)

☒ CC (Header)

☒ Reply-To (Header)

▸ Group

Enhanced Support for Microsoft Exchange 5.5

AsyncOS includes a configuration option to provide support for Microsoft Exchange 5.5. If you use a later version of Microsoft Exchange, you do not need to enable this option. When configuring an LDAP server, you can elect to enable Microsoft Exchange 5.5 support by answering “y” when prompted in the `ldapconfig` → `edit` → `server` → `compatibility` subcommand (this is only available via the CLI):

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

```
Current LDAP server configurations:
```



```
1. PublicLDAP: (ldapexample.com:389)
```

```
Choose the operation you want to perform:
```

- NEW - Create a new server configuration.
- EDIT - Modify a server configuration.
- DELETE - Remove a server configuration.

```
[> edit
```

```
Enter the name or number of the server configuration you wish to edit.
```

```
[> 1
```

```
Name: PublicLDAP
```

```
Hostname: ldapexample.com Port 389
```

```
Authentication Type: anonymous
```

```
Base: dc=ldapexample,dc=com
```

```
Choose the operation you want to perform:
```

- SERVER - Change the server for the query.
- LDAPACCEPT - Configure whether a recipient address should be accepted or bounced/dropped.
- LDAPROUTING - Configure message routing.
- MASQUERADE - Configure domain masquerading.
- LDAPGROUP - Configure whether a sender or recipient is in a specified group.
- SMTPAUTH - Configure SMTP authentication.

```
[> server
```

```
Name: PublicLDAP
```

```
Hostname: ldapexample.com Port 389
```

```
Authentication Type: anonymous
```

Base: dc=ldapexample,dc=com

Microsoft Exchange 5.5 Compatibility Mode: Disabled

Choose the operation you want to perform:

- NAME - Change the name of this configuration.
- HOSTNAME - Change the hostname used for this query.
- PORT - Configure the port.
- AUTHTYPE - Choose the authentication type.
- BASE - Configure the query base.
- COMPATIBILITY - Set LDAP protocol compatibility options.

[> compatibility

Would you like to enable Microsoft Exchange 5.5 LDAP compatibility mode? (This is not recommended for versions of Microsoft Exchange later than 5.5, or other LDAP servers.)

[N]> y

Do you want to configure advanced LDAP compatibility settings? (Typically not required)

[N]>

Name: PublicLDAP

Hostname: ldapexample.com Port 389

Authentication Type: anonymous

Base: dc=ldapexample,dc=com

Microsoft Exchange 5.5 Compatibility Mode: Enabled (attribute "objectClass")

Choose the operation you want to perform:

- NAME - Change the name of this configuration.
- HOSTNAME - Change the hostname used for this query.
- PORT - Configure the port.
- AUTHTYPE - Choose the authentication type.

```
- BASE - Configure the query base.  
  
- COMPATIBILITY - Set LDAP protocol compatibility options.  
  
[]>
```

Working with LDAP Queries

You create an entry in the LDAP server profile for each type of LDAP query you want to perform. When you create LDAP queries, you must enter the query syntax for your LDAP server. Please note that the queries you construct should be tailored and specific to your particular implementation of LDAP directory services, particularly if you have extended your directory with new object classes and attributes to accommodate the unique needs of your directory.

Related Topics

- [Types of LDAP Queries, page 25-12](#)
- [Base Distinguishing Name \(DN\), page 25-13](#)
- [LDAP Query Syntax, page 25-13](#)
- [Secure LDAP \(SSL\), page 25-14](#)
- [Routing Queries, page 25-14](#)
- [Allowing Clients to Bind to the LDAP Server Anonymously, page 25-14](#)
- [Testing LDAP Queries, page 25-17](#)
- [Troubleshooting Connections to LDAP Servers, page 25-18](#)

Types of LDAP Queries

- **Acceptance queries.** For more information, see [Using Acceptance Queries For Recipient Validation, page 25-19](#).
- **Routing queries.** For more information, see [Using Routing Queries to Send Mail to Multiple Target Addresses, page 25-20](#).
- **Certificate Authentication queries.** For more information, see [Checking the Validity of a Client Certificate, page 26-51](#).
- **Masquerading queries.** For more information, see [Using Masquerading Queries to Rewrite the Envelope Sender, page 25-21](#).
- **Group queries.** For more information, see [Using Group LDAP Queries to Determine if a Recipient is a Group Member, page 25-23](#).
- **Domain-based queries.** For more information, see [Using Domain-based Queries to Route to a Particular Domain, page 25-26](#).
- **Chain queries.** For more information, see [Using Chain Queries to Perform a Series of LDAP Queries, page 25-28](#).

You can also configure queries for the following purposes:

- **Directory harvest prevention.** For more information, see [Understanding LDAP Queries, page 25-2](#).

- **SMTP authentication.** For more information, see [Configuring AsyncOS for SMTP Authentication, page 25-32](#).
- **External authentication.** For more information, see [Configuring External LDAP Authentication for Users, page 25-40](#).
- **Spam quarantine end-user authentication query.** For more information, see [Authenticating End-Users of the Spam Quarantine, page 25-43](#).
- **Spam quarantine alias consolidation query.** For more information, see [Spam Quarantine Alias Consolidation Queries, page 25-44](#).

The search queries you specify are available to all listeners you configure on the system.

Base Distinguishing Name (DN)

The root level of the directory is called the base. The name of the base is the DN (distinguishing name). The base DN format for Active Directory (and the standard as per RFC 2247) has the DNS domain translated into domain components (dc=). For example, example.com's base DN would be: dc=example, dc=com. Note that each portion of the DNS name is represented in order. This may or may not reflect the LDAP settings for your configuration.

If your directory contains multiple domains you may find it inconvenient to enter a single BASE for your queries. In this case, when configuring the LDAP server settings, set the base to NONE. This will, however, make your searches inefficient.

LDAP Query Syntax

Spaces are allowed in LDAP paths, and they do not need to be quoted. The CN and DC syntax is not case-sensitive.

```
Cn=First Last,oU=user,dc=domain,DC=COM
```

The variable names you enter for queries are *case-sensitive* and must match your LDAP implementation in order to work correctly. For example, entering **mailLocalAddress** at a prompt performs a different query than entering **maillocaladdress**.

Related Topics

- [Tokens:, page 25-13](#)

Tokens:

You can use the following tokens in your LDAP queries:

- {a} username@domainname
- {d} domainname
- {dn} distinguished name
- {g} groupname
- {u} username
- {f} MAIL FROM: address



Note The {f} token is valid in acceptance queries only.

For example, you might use the following query to accept mail for an Active Directory LDAP server:
(l(mail={a})(proxyAddresses=smtp:{a}))



Note

Cisco Systems strongly recommends using the Test feature of the LDAP page (or the `test` subcommand of the `ldapconfig` command) to test all queries you construct and ensure that expected results are returned *before* you enable LDAP functionality on a listener. See [Testing LDAP Queries, page 25-17](#) for more information.

Secure LDAP (SSL)

You can use instruct AsyncOS to use SSL when communicating with the LDAP server. If you configure your LDAP server profile to use SSL:

- AsyncOS will use the LDAPS certificate configured via `certconfig` in the CLI (see [Creating a Self-Signed Certificate using the GUI, page 23-3](#)).

You may have to configure your LDAP server to support using the LDAPS certificate.

- If an LDAPS certificate has not been configured, AsyncOS will use the demo certificate.

Routing Queries

There is no recursion limit for LDAP routing queries; the routing is completely data driven. However, AsyncOS does check for circular reference data to prevent the routing from looping infinitely.

Allowing Clients to Bind to the LDAP Server Anonymously

You may need to configure your LDAP directory server to allow for anonymous queries. (That is, clients can bind to the server anonymously and perform queries.) For specific instructions on configuring Active Directory to allow anonymous queries, see the “Microsoft Knowledge Base Article - 320528” at the following URL:

<http://support.microsoft.com/default.aspx?scid=kb%3Ben-us%3B320528>

Alternately, you can configure one “user” dedicated solely for the purposes of authenticating and performing queries instead of opening up your LDAP directory server for anonymous queries from any client.

A summary of the steps is included here, specifically:

- How to set up Microsoft Exchange 2000 server to allow “anonymous” authentication.
- How to set up Microsoft Exchange 2000 server to allow “anonymous bind.”
- How to set up AsyncOS to retrieve LDAP data from a Microsoft Exchange 2000 server using both “anonymous bind” and “anonymous” authentication.

Specific permissions must be made to a Microsoft Exchange 2000 server in order to allow “anonymous” or “anonymous bind” authentication for the purpose of querying user email addresses. This can be very useful when an LDAP query is used to determine the validity of an incoming email message to the SMTP gateway.

Related Topics

- [Anonymous Authentication Setup, page 25-15](#)
- [Anonymous Bind Setup for Active Directory, page 25-16](#)
- [Notes for Active Directory Implementations, page 25-17](#)

Anonymous Authentication Setup

The following setup instructions allow you to make specific data available to unauthenticated queries of Active Directory and Exchange 2000 servers in the Microsoft Windows Active Directory. If you wish to allow “anonymous bind” to the Active Directory, see [Anonymous Bind Setup for Active Directory, page 25-16](#).

Procedure

Step 1 Determine required Active Directory permissions.

Using the ADSI Edit snap-in or the LDP utility, you must modify the permissions to the attributes of the following Active Directory objects:

- The root of the domain naming context for the domain against which you want to make queries.
- All OU and CN objects that contain users against which you wish to query email information.

The following table shows the required permissions to be applied to all of the needed containers.

User Object	Permissions	Inheritance	Permission Type
Everyone	List Contents	Container Objects	Object
Everyone	List Contents	Organizational Unit Objects	Object
Everyone	Read Public Information	User Objects	Property
Everyone	Read Phone and Mail Options	User Objects	Property

Step 2 Set Active Directory Permissions

- Open ADSIEdit from the Windows 2000 Support Tools.
- Locate the **Domain Naming Context** folder. This folder has the LDAP path of your domain.
- Right click the **Domain Naming Context** folder, and then click **Properties**.
- Click **Security**.
- Click **Advanced**.
- Click **Add**.
- Click the **User Object** Everyone, and then click **OK**.
- Click the **Permission Type** tab.
- Click **Inheritance** from the **Apply onto** box.

- Click to select the Allow check box for the **Permission** permission.

Step 3 Configure the Cisco Messaging Gateway

Use `ldapconfig` on the Command Line Interface (CLI) to create an LDAP server entry with the following information.

- Hostname of an Active Directory or Exchange server
- Port 3268
- Base DN matching the root naming context of the domain
- Authentication type Anonymous

Anonymous Bind Setup for Active Directory

The following setup instructions allow you to make specific data available to anonymous bind queries of Active Directory and Exchange 2000 servers in the Microsoft Windows Active Directory. Anonymous bind of an Active Directory server will send the username `anonymous` with a blank password.



Note

If a password is sent to an Active Directory server while attempting anonymous bind, authentication may fail.

Procedure

Step 1 Determine required Active Directory permissions.

Using the ADSI Edit snap-in or the LDP utility, you must modify the permissions to the attributes of the following Active Directory objects.

- The root of the domain naming context for the domain against which you want to make queries.
- All OU and CN objects that contain users against which you wish to query email information.

The following table shows the required permissions to be applied to all of the needed containers.

User Object	Permissions	Inheritance	Permission Type
ANONYMOUS LOGON	List Contents	Container Objects	Object
ANONYMOUS LOGON	List Contents	Organizational Unit Objects	Object
ANONYMOUS LOGON	Read Public Information	User Objects	Property
ANONYMOUS LOGON	Read Phone and Mail Options	User Objects	Property

Step 2 Set Active Directory Permissions

- Open ADSIEdit from the Windows 2000 Support Tools.
- Locate the **Domain Naming Context** folder. This folder has the LDAP path of your domain.
- Right click the **Domain Naming Context** folder, and then click **Properties**.
- Click **Security**.
- Click **Advanced**.
- Click **Add**.

- Click the **User Object** ANONYMOUS LOGON, and then click **OK**.
- Click the **Permission Type** tab.
- Click **Inheritance** from the **Apply** onto box.
- Click to select the **Allow** check box for the **Permission** permission.

Step 3 Configure the Cisco Messaging Gateway

Use the System Administration > LDAP page (or `ldapconfig` in the CLI) to create an LDAP server entry with the following information.

- Hostname of an Active Directory or Exchange server
- Port 3268
- Base DN matching the root naming context of the domain
- Authentication type password based using `cn=anonymous` as the user with a blank password

Notes for Active Directory Implementations

- Active Directory servers accept LDAP connections on ports 3268 and 389. The default port for accessing the global catalog is port 3268.
- Active Directory servers accept LDAPS connections on ports 636 and 3269. Microsoft supports LDAPS on Windows Server 2003 and higher.
- The appliance should connect to a domain controller that is also a global catalog so that you can perform queries to different bases using the same server.
- Within Active Directory, you may need to grant read permissions to the group “Everyone” to directory objects to yield successful queries. This includes the root of the domain naming context.
- Generally, the value of the `mail` attribute entry in many Active Directory implementations has a matching value “ProxyAddresses” attribute entry.
- Microsoft Exchange environments that are aware of each other within the infrastructure can usually route mail between each other without involving a route back to the originating MTA.

Testing LDAP Queries

Use the Test Query button on the Add/Edit LDAP Server Profile page (or the `test` subcommand in the CLI) of each query type to test the query to the LDAP server you configured. In addition to displaying the result, AsyncOS also displays the details on each stage of the query connection test. You can test each of the query types.

The `ldaptest` command is available as a batch command, for example:

```
ldaptest LDAP.ldapaccept foo@ironport.com
```


If you entered multiple hosts in the Host Name field of the LDAP server attributes, the appliance tests the query on each LDAP server.

Table 25-1 *Testing LDAP Queries*

Query type	If a recipient matches (PASS)...	If a recipient does not match (FAIL)...
Recipient Acceptance (Accept, ldapaccept)	Accept the message.	Invalid Recipient: Conversation or delayed bounce or drop the message per listener settings. DHAP: Drop.
Routing (Routing, ldaprouting)	Route based on the query settings.	Continue processing the message.
Masquerade (Masquerade, masquerade)	Alter the headers with the variable mappings defined by the query.	Continue processing the message.
Group Membership (Group, ldapgroup)	Return “true” for message filter rules.	Return “false” for message filter rules.
SMTP Auth (SMTP Authentication, smtpauth)	A password is returned from the LDAP server and is used for authentication; SMTP Authentication occurs.	No password match can occur; SMTP Authentication attempts fail.
External Authentication (externalauth)	Individually returns a “match positive” for the bind, the user record, and the user’s group membership.	Individually returns a “match negative” for the bind, the user record, and the user’s group membership.
Spam Quarantine End-User Authentication (isqauth)	Returns a “match positive” for the end-user account.	No password match can occur; End-User Authentication attempts fail.
Spam Quarantine Alias Consolidation (isqalias)	Returns the email address that the consolidated spam notifications will be sent to.	No consolidation of spam notifications can occur.



Note

The variable names you enter for queries are *case-sensitive* and must match your LDAP implementation in order to work correctly. For example, entering **mailLocalAddress** at a prompt performs a different query than entering **maillocaladdress**. Cisco Systems strongly recommends using the `test` subcommand of the `ldapconfig` command to test all queries you construct and ensure the proper results are returned.

Troubleshooting Connections to LDAP Servers

If the LDAP server is unreachable by the appliance, one of the following errors will be shown:

- Error: LDAP authentication failed: <LDAP Error "invalidCredentials" [0x31]>
- Error: Server unreachable: unable to connect
- Error: Server unreachable: DNS lookup failure

Note that a server may be unreachable because the wrong port was entered in the server configuration, or the port is not opened in the firewall. LDAP servers typically communicate over port 3268 or 389. Active Directory uses port 3268 to access the global catalog used in multi-server environments (See the “Firewall Information” appendix for more information.) In AsyncOS 4.0, the ability to communicate to the LDAP server via SSL (usually over port 636) was added. For more information, see [Secure LDAP \(SSL\)](#), page 25-14.

A server may also be unreachable because the hostname you entered cannot be resolved.

You can use the **Test Server(s)** on the Add/Edit LDAP Server Profile page (or the `test` subcommand of the `ldapconfig` command in the CLI) to test the connection to the LDAP server. For more information, see [Testing LDAP Servers](#), page 25-6.

If the LDAP server is unreachable:

- If LDAP Accept or Masquerading or Routing is enabled on the work queue, mail will remain within the work queue.
- If LDAP Accept is not enabled but other queries (group policy checks, etc.) are used in filters, the filters evaluate to false.

Using Acceptance Queries For Recipient Validation

You can use your existing LDAP infrastructure to define how the recipient email address of incoming messages (on an public listener) should be handled. Changes to user data in your directories are updated the next time the appliance queries the directory server. You can specify the size of the caches and the amount of time the appliance stores the data it retrieves.



Note

You may wish to bypass LDAP acceptance queries for special recipients (such as `administrator@example.com`). You can configure this setting from the Recipient Access Table (RAT). For information about configuring this setting, see the “Configuring the Gateway to Receive Email” chapter.

Related Topics

- [Sample Acceptance Queries](#), page 25-19
- [Configuring Acceptance Queries for Lotus Notes](#), page 25-20

Sample Acceptance Queries

Table 25-2 shows sample acceptance queries.

Table 25-2 Example LDAP Query Strings for Common LDAP Implementations: Acceptance

Query for:	Recipient validation
OpenLDAP	<pre>(mailLocalAddress={a}) (mail={a}) (mailAlternateAddress={a})</pre>
Microsoft Active Directory Address Book Microsoft Exchange	<pre>((mail={a}) (proxyAddresses=smtp:{a}))</pre>

Table 25-2 Example LDAP Query Strings for Common LDAP Implementations: Acceptance

Query for:	Recipient validation
SunONE Directory Server	<pre>(mail={a}) (mailAlternateAddress={a}) (mailEquivalentAddress={a}) (mailForwardingAddress={a}) (mailRoutingAddress={a})</pre>
Lotus Notes Lotus Domino	<pre>(((mail={a}) (uid={u})) (cn={u})) ((ShortName={u}) (InternetAddress={a}) (FullNa me={u}))</pre>

You can also validate on the username (Left Hand Side). This is useful if your directory does not contain all the domains you accept mail for. Set the Accept query to `(uid={u})`.

Configuring Acceptance Queries for Lotus Notes

Note that there is a potential complication with LDAPACCEPT and Lotus Notes. If Notes LDAP contains a person with attributes like these:

```
mail=juser@example.com
```

```
cn=Joe User
```

```
uid=juser
```

```
cn=123456
```

```
location=New Jersey
```

Lotus accepts email for this person for various different forms of email addresses, other than what is specified, such as “Joe_User@example.com” — which do not exist in the LDAP directory. So AsyncOS may not be able to find all of the valid user email addresses for that user.

One possible solution is to try to publish the other forms of addresses. Please contact your Lotus Notes administrator for more details.

Using Routing Queries to Send Mail to Multiple Target Addresses

AsyncOS supports alias expansion (LDAP routing with multiple target addresses). AsyncOS replaces the original email message with a new, separate message for each alias target (for example, `recipient@yoursite.com` might be replaced with new separate messages to `newrecipient1@hotmail.com` and `recipient2@internal.yourcompany.com`, etc.). Routing queries are sometimes known as aliasing queries on other mail processing systems.

Related Topics

- [Sample Routing Queries, page 25-21](#)

Sample Routing Queries

Table 25-3 *Example LDAP Query Strings for Common LDAP Implementations: Routing*

Query for:	Route to another mailhost
OpenLDAP	(mailLocalAddress={a})
Microsoft Active Directory Address Book	May not be applicable ^a
Microsoft Exchange	
SunONE Directory Server	(mail={a}) (mailForwardingAddress={a}) (mailEquivalentAddress={a}) (mailRoutingAddress={a}) (otherMailbox={a}) (rfc822Mailbox={a})

a.Active Directory implementations can have multiple entries for the proxyAddresses attribute, but because AD formats this attribute value as smtp:user@domain.com, that data cannot be used for LDAP routing/alias expansion. Each target address must be in a separate attribute:value pair. Microsoft Exchange environments that are aware of each other within the infrastructure can usually route mail between each other without involving a route back to the originating MTA.

Related Topics

- [Routing: MAILHOST and MAILROUTINGADDRESS, page 25-21](#)

Routing: MAILHOST and MAILROUTINGADDRESS

For Routing queries, the value of MAILHOST cannot be an IP address; it must be a resolvable hostname. This usually requires the use of an Internal DNSconfig.

MAILHOST is optional for the routing query. MAILROUTINGADDRESS is mandatory if MAILHOST is not set.

Using Masquerading Queries to Rewrite the Envelope Sender

Masquerading is a feature that rewrites the Envelope Sender (also known as the sender, or MAIL FROM) and the To:, From:, and/or CC: headers on email based on queries you construct. A typical example implementation of this feature is “Virtual Domains,” which allows you to host multiple domains from a single site. Another typical implementation is “hiding” your network infrastructure by “stripping” the subdomains from strings in email headers.

Related Topics

- [Sample Masquerading Queries, page 25-22](#)
- [Masquerading “Friendly Names”, page 25-22](#)

Sample Masquerading Queries

Table 25-4 Example LDAP Query Strings for Common LDAP Implementation: Masquerading

Query for:	Masquerade
OpenLDAP	(mailRoutingAddress={a})
Microsoft Active Directory Address Book	(proxyaddresses=smtp:{a})
SunONE Directory Server	(mail={a}) (mailAlternateAddress={a}) (mailEquivalentAddress={a}) (mailForwardingAddress={a}) (mailRoutingAddress={a})

Masquerading “Friendly Names”

In some user environments, an LDAP directory server schema may store a “friendly name” in addition to a mail routing address or a local mail address. AsyncOS allows you to masquerade Envelope Senders (for outgoing mail) and message headers (for incoming mail, such as To:, Reply To:, From: or CC:) with this “friendly address” — even if the friendly address contains special characters that are not normally permitted in a valid email address (for example, quotation marks, spaces, and commas).

When using masquerading of headers via an LDAP query, you now have the option to configure whether to replace the entire friendly email string with the results from the LDAP server. Note that even with this behavior enabled, only the user@domain portion will be used for the Envelope Sender (the friendly name is illegal).

As with the normal LDAP masquerading, if empty results (zero length or entire white space) are returned from the LDAP query, no masquerading occurs.

To enable this feature, answer “y” to the following question when configuring an LDAP-based masquerading query for a listener (LDAP page or `ldapconfig` command):

```
Do you want the results of the returned attribute to replace the entire
friendly portion of the original recipient? [N]
```

For example, consider the following example LDAP entry:

Attribute	Value
mailRoutingAddress	admin\@example.com
mailLocalAddress	joe.smith\@example.com
mailFriendlyAddress	“Administrator for example.com,” <joe.smith\@example.com>

If this feature is enabled, an LDAP query of (mailRoutingAddress={a}) and a masquerading attribute of (mailLocalAddress) would result in the following substitutions:

Original Address (From, To, CC, Reply-to)	Masqueraded Headers	Masqueraded Envelope Sender
admin@example.com	From: “Administrator for example.com,” <joe.smith@example.com>	MAIL FROM: <joe.smith@example.com>

Using Group LDAP Queries to Determine if a Recipient is a Group Member

You can define a query to your LDAP servers to determine if a recipient is a member of a group as defined by your LDAP directory.

Procedure

-
Step 1	Create a message filter that uses a <code>rcpt-to-group</code> or <code>mail-from-group</code> rule to act upon the message.
Step 2	Then, use the System Administration > LDAP page (or the <code>ldapconfig</code> command) to define the LDAP server for the appliance to bind to and configure a query for a group membership.
Step 3	Use the Network > Listeners page (or the <code>listenerconfig -> edit -> ldapgroup</code> subcommand) to enable the group query for the listener.
-

Related Topics

- [Sample Group Queries, page 25-23](#)
- [Configuring a Group Query, page 25-23](#)
- [Example: Using a Group Query to Skip Spam and Virus Checking, page 25-25](#)

Sample Group Queries

Table 25-5 *Example LDAP Query Strings for Common LDAP Implementation: Group*

Query for:	Group
OpenLDAP	OpenLDAP does not support the <code>memberOf</code> attribute by default. Your LDAP Administrator may add this attribute or a similar attribute to the schema.
Microsoft Active Directory	<code>(&(memberOf={g})(proxyAddresses=smtp:{a}))</code>
SunONE Directory Server	<code>(&(memberOf={g})(mailLocalAddress={a}))</code>

For example, suppose that your LDAP directory classifies members of the “Marketing” group as `ou=Marketing`. You can use this classification to treat messages sent to or from members of this group in a special way. Step 1 creates a message filter to act upon the message, and Steps 2 and 3 enable the LDAP lookup mechanism.

Configuring a Group Query

In the following example, mail from members of the Marketing group (as defined by the LDAP group “Marketing”) will be delivered to the alternate delivery host `marketingfolks.example.com`.

Procedure

Step 1 First, a message filter is created to act upon messages that match positively for group membership. In this example, a filter is created that uses the `mail-from-group` rule. All messages whose Envelope Sender is found to be in the LDAP group “marketing-group1” will be delivered with an alternate delivery host (the filters `alt-mailhost` action).

The group membership field variable (`groupName`) will be defined in step 2. The group attribute “`groupName`” is defined with the value `marketing-group1`.

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

Choose the operation you want to perform:

- NEW - Create a new filter.
- IMPORT - Import a filter script from a file.

```
[> new
```

Enter filter script. Enter '.' on its own line to end.

MarketingGroupfilter:

```
if (mail-from-group == "marketing-group1") {
    alt-mailhost ('marketingfolks.example.com');}
.
```

```
1 filters added.
```

Choose the operation you want to perform:

- NEW - Create a new filter.
- DELETE - Remove a filter.
- IMPORT - Import a filter script from a file.
- EXPORT - Export filters to a file
- MOVE - Move a filter to a different position.
- SET - Set a filter attribute.
- LIST - List the filters.
- DETAIL - Get detailed information on the filters.

- LOGCONFIG - Configure log subscriptions used by filters.
 - ROLLOVERNOW - Roll over a filter log file.
- []>

For more information on the `mail-from-group` and `rcpt-to-group` message filter rules, see [Message Filter Rules, page 9-2](#).

Step 2 Next, the Add LDAP Server Profile page is used to define an LDAP server for the appliance to bind to, and an initial query for a group membership is configured.

Step 3 Next, the public listener “InboundMail” is updated to use LDAP queries for group routing. The Edit Listener page is used to enable the LDAP query specified above.

As a result of this query, messages accepted by the listener trigger a query to the LDAP server to determine group membership. The `PublicLDAP2.group` query was defined previously via the System Administration > LDAP page.

Figure 25-6 Specifying a Group Query on a Listener
Edit Listener

Listener Settings	
Name:	IncomingMail
Type of Listener:	Public
Interface:	Data 1 TCP Port: 25
Bounce Profile:	Default
Disclaimer Above:	None Disclaimer text will be applied above the message body.
Disclaimer Below:	None Disclaimer text will be applied below the message body.
SMTP Authentication Profile:	None
Certificate:	test
▶ SMTP Address Parsing Options:	Optional settings for controlling parsing in SMTP "MAIL FROM" and "RCPT TO" fields.
▶ Advanced:	Optional settings for customizing the behavior of the Listener
▼ LDAP Queries:	▶ Accept ▶ Routing ▶ Masquerade ▼ Group

Note that in this example, a commit must be issued for the changes to take effect.

Example: Using a Group Query to Skip Spam and Virus Checking

Because message filters occurs early in the pipeline, you can use a group query to skip virus and spam checking for specified groups. For example, you want your IT group to receive all messages and to skip spam and virus checking. In your LDAP record, you create a group entry that uses the DN as the group name. The group name consists of the following DN entry:

```
cn=IT, ou=groups, o=sample.com
```

You create an LDAP server profile with the following group query:

```
(&(memberOf={g})(proxyAddresses=smtp:{a}))
```


You then enable this query on a listener so that when a message is received by the listener, the group query is triggered.

To skip virus and spam filtering for members of the IT group, you create the following message filter to check incoming messages against LDAP groups.

```
[> - NEW - Create a new filter.

- IMPORT - Import a filter script from a file.

[> new

Enter filter script. Enter '.' on its own line to end.

IT_Group_Filter:

if (rcpt-to-group == "cn=IT, ou=groups, o=sample.com"){

skip-spamcheck();

skip-viruscheck();

deliver();

}

.

1 filters added.
```


Note

The rcpt-to-group in this message filter reflects the DN entered as the group name: cn=IT, ou=groups, o=sample.com. Verify that you use the correct group name in the message filter to ensure that your filter matches the name in your LDAP directory.

Messages accepted by the listener trigger a query to the LDAP server to determine group membership. If the message recipient is a member of the IT group, the message filter skips both virus and spam checking and delivers the message to the recipient. To enable the filter to check the results of the LDAP query, you must create the LDAP query on the LDAP server and enable the LDAP query on a listener.

Using Domain-based Queries to Route to a Particular Domain

Domain-based queries are LDAP queries grouped by type, associated with a domain, and assigned to a particular listener. You might want to use domain-based queries if you have different LDAP servers associated with different domains but you want to run queries for all your LDAP servers on the same listener. For example, the company “MyCompany” purchases company “HisCompany” and company “HerCompany” MyCompany maintains its domain, MyCompany.example.com as well as domains for HisCompany.example.com and HerCompany.example.com, and it maintains a different LDAP server for employees associated with each domain. To accept mail for all three of these domains, MyCompany

creates domain-based queries. This allows MyCompany.example.com to accept emails for Mycompany.example.com, HisCompany.example.com, and HerCompany.example.com on the same listener.

Procedure

-
- Step 1** Create a server profile for each of the domains you want to use in the domain-based queries. For each of the server profiles, configure the queries you want to use for a domain-based query (acceptance, routing, etc.). For more information, see [Creating LDAP Server Profiles to Store Information About the LDAP Server, page 25-5](#).
- Step 2** Create the domain-based query. When you create the domain-based query, you select queries from each server profile, and enable the appliance to determine which query to run based on the domain in the Envelope To field. For more information about creating the query, see [Creating a Domain-Based Query, page 25-27](#).
- Step 3** Enable the domain-based query on the public or private listener. For more information about configuring listeners, see the “Configuring the Gateway to Receive Mail” chapter.



Note

You can also enable domain-based queries for LDAP end-user access or spam notifications for the Spam Quarantine. For more information, see the Spam Quarantine chapter.

Related Topics

- [Creating a Domain-Based Query, page 25-27](#)

Creating a Domain-Based Query

You create a domain-based query from the System Administration > LDAP > LDAP Server Profiles page.

Procedure

-
- Step 1** From the LDAP Server Profiles page, click **Advanced**.
- Step 2** Click **Add Domain Assignments**.
- Step 3** Enter a name for the domain-based query.
- Step 4** Select the query type.



Note

When you create domain-based queries, you cannot select different types of queries. Once you select a query type, the appliance populates the query field with queries of that type from the available server profiles.

- Step 5** In the Domain Assignments field, enter a domain.
- Step 6** Select a query to associate with the domain.
- Step 7** Continue to add rows until you have added all the domains to your query.

- Step 8** You can enter a default query to run if all other queries fail. If you do not want to enter a default query, select **None**.
- Step 9** Test the query by clicking the **Test Query** button and entering a user login and password or an email address to test in the Test Parameters fields. The results appear in the Connection Status field.
- Step 10** Optionally, if you use the {f} token in an acceptance query, you can add an envelope sender address to the test query.



Note Once you create the domain-based query, you need to associate it with a public or private listener.

- Step 11** Submit and commit your changes.

Using Chain Queries to Perform a Series of LDAP Queries

A chain query is a series of LDAP queries that the appliance attempts to run in succession. The appliance attempts to run each query in the “chain” until the LDAP server returns a positive response (or the final query in the “chain” returns a negative response or fails). Chain queries can be useful if entries in your LDAP directory use different attributes to store similar (or the same) values. For example, you might have used the attributes `maillocaladdress` and `mail` to store user email addresses. To ensure that your queries run against both these attributes, you can use chain queries.

Procedure

- Step 1** Create server profiles for each of the queries you want to use in the chain queries. For each of the server profiles, configure the queries you want to use for a chain query. For more information, see [Creating LDAP Server Profiles to Store Information About the LDAP Server, page 25-5](#).
- Step 2** Create the chain query. For more information, see [Creating a Chain Query, page 25-28](#).
- Step 3** Enable the chain query on the public or private listener. For more information about configuring listeners, see the “Configuring the Gateway to Receive Mail” chapter.



Note You can also enable domain-based queries for LDAP end-user access or spam notifications for the Spam Quarantine. For more information, see the Spam Quarantine chapter.

Related Topics

- [Creating a Chain Query, page 25-28](#)

Creating a Chain Query

You create a chain query from the System Administration > LDAP > LDAP Server Profiles page.

Procedure

-
- Step 1** From the LDAP Server Profiles page, click **Advanced**.
- Step 2** Click **Add Chain Query**.
- Step 3** Add a name for the chain query.
- Step 4** Select the query type.
- When you create chain queries, you cannot select different types of queries. Once you select a query type, the appliance populates the query field with queries of that type from available server profiles.
- Step 5** Select a query to add to the chain query.
- The appliance runs the queries in the order you configure them. Therefore, if you add multiple queries to the chain query, you might want to order the queries so that more specific queries are followed by more general queries.
- Step 6** Test the query by clicking the **Test Query** button and entering a user login and password or an email address to test in the Test Parameters fields. The results appear in the Connection Status field.
- Step 7** Optionally, if you use the {f} token in an acceptance query, you can add an envelope sender address to the test query.
- 

Note Once you create the chain query, you need to associate it with a public or private listener.
-
- Step 8** Submit and commit your changes.
-

Using LDAP For Directory Harvest Attack Prevention

Directory Harvest Attacks occur when a malicious sender attempts to send messages to recipients with common names, and the email gateway responds by verifying that a recipient has a valid mailbox at that location. When performed on a large scale, malicious senders can determine who to send mail to by “harvesting” these valid addresses for spamming.

The Email Security appliance can detect and prevent Directory Harvest Attack (DHA) when using LDAP acceptance validation queries. You can configure LDAP acceptance to prevent directory harvest attacks within the SMTP conversation or within the work queue.

Related Topics

- [Directory Harvest Attack Prevention within the SMTP Conversation, page 25-29](#)
- [Directory Harvest Attack Prevention within the Work Queue, page 25-31](#)

Directory Harvest Attack Prevention within the SMTP Conversation

You can prevent DHAs by entering only domains in the Recipient Access Table (RAT), and performing the LDAP acceptance validation in the SMTP conversation.

To drop messages during the SMTP conversation, configure an LDAP server profile for LDAP acceptance. Then, configure the listener to perform an LDAP accept query during the SMTP conversation.

Figure 25-7 Configuring the Acceptance Query in the SMTP Conversation

LDAP Queries: Accept

Accept Query: redfish.accept

Work Queue

Non-Matching Recipients: Bounce

SMTP Conversation

If the LDAP server is unreachable:

Allow Mail in

Return error code:

Code: 451

Text: Temporary recipient validation er

Routing

Masquerade

Group

Once you configure LDAP acceptance queries for the listener, you must configure DHAP settings in the mail flow policy associated with the listener.

Figure 25-8 Configuring the Mail Flow Policy to Drop Connections in the SMTP Conversation

Mail Flow Limits		
Rate Limiting:	Max. Recipients Per Hour:	Unlimited
	Max. Recipients Per Hour Code:	452
	Max. Recipients Per Hour Text:	Too many recipients received this hour
Flow Control:	Use SenderBase for Flow Control:	On Off
	Group by Similarity of IP Addresses:	This Feature can only be used if Senderbase Flow Control is off. Off
		(significant bits 0-32)
Directory Harvest Attack Prevention (DHAP):	Max. Invalid Recipients Per Hour:	Unlimited 5
	Drop Connection if DHAP threshold is Reached within an SMTP Conversation:	On Off
	Max. Invalid Recipients Per Hour Code:	550
	Max. Invalid Recipients Per Hour Text:	Too many invalid recip

In the mail flow policy associated with the listener, configure the following Directory Harvest Attack Prevention settings:

- **Max. 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. For example, you configure the threshold as five, and the counter detects two RAT rejections and three dropped messages to invalid LDAP recipients. At this point, the appliance determines that the threshold is reached, and the connection is dropped. By default, the maximum number of recipients per hour for a public listener is 25. For a private listener, the maximum number of recipients per hour is unlimited by default. Setting it to “Unlimited” means that DHAP is not enabled for that mail flow policy.
- **Drop Connection if DHAP Threshold is reached within an SMTP conversation.** Configure the appliance to drop the connection if the Directory Harvest Attack Prevention threshold is reached.
- **Max. Recipients Per Hour Code.** Specify the code to use when dropping connections. The default code is 550.
- **Max. Recipients Per Hour Text.** Specify the text to use for dropped connections. The default text is “Too many invalid recipients.”

If the threshold is reached, the Envelope Sender of the message does not receive a bounce message when a recipient is invalid.

Directory Harvest Attack Prevention within the Work Queue

You can prevent most DHAs by entering only domains in the Recipient Access Table (RAT), and performing the LDAP acceptance validation within the work queue. This technique prevents the malicious senders from knowing if the recipient is valid during the SMTP conversation. (When acceptance queries are configured, the system accepts the message and performs the LDAP acceptance validation within the work queue.) However, the Envelope Sender of the message will still receive a bounce message if a recipient is not valid.


Related Topics

- [Configuring Directory Harvest Prevention in the Work Queue, page 25-31](#)

Configuring Directory Harvest Prevention in the Work Queue

To prevent Directory Harvest Attacks, you first configure an LDAP server profile, and enable LDAP Accept. Once you have enabled LDAP acceptance queries, configure the listener to use the accept query, and to bounce mail for non-matching recipients:

Figure 25-9 *Configuring the Acceptance Query to Bounce Messages for Non-Matching Recipients*



Next, configure the Mail Flow Policy to define the number of invalid recipient addresses the system will allow per sending IP address for a specific period of time. When this number is exceeded, the system will identify this condition as a DHA and send an alert message. The alert message will contain the following information:

```
LDAP: Potential Directory Harvest Attack from host=('IP-address', 'domain_name'),
dhap_limit=n, sender_group=sender_group,
```

```
listener=listener_name, reverse_dns=(reverse_IP_address, 'domain_name', 1),
sender=envelope_sender, rcpt=envelope_recipients
```

The system will bounce the messages up to the threshold you specified in the mail flow policy and then it will silently accept and drop the rest, thereby informing legitimate senders that an address is bad, but preventing malicious senders from determining which receipts are accepted.

This invalid recipients counter functions similarly to the way Rate Limiting is currently available in AsyncOS: you enable the feature and define the limit as part of the mail flow policy in a public listener's HAT (including the default mail flow policy for the HAT).

For example, you are prompted with these questions when creating or editing a mail flow policy in a public listener's HAT in the CLI — the `listenerconfig -> edit -> hostaccess -> default | new` commands:


```
Do you want to enable Directory Harvest Attack Prevention per host? [Y]> y
```

Enter the maximum number of invalid recipients per hour from a remote host.

[25]>

This feature is also displayed when editing any mail flow policy in the GUI, providing that LDAP queries have been configured on the corresponding listener:

Figure 25-10 DHAP Prevention Feature in GUI



Entering a number of invalid recipients per hour enables DHAP for that mail flow policy. By default, 25 invalid recipients per hour are allowed for public listeners. For private listeners, the maximum invalid recipients per hour is unlimited by default. Setting it to “Unlimited” means that DHAP is not enabled for that mail flow policy.

Configuring AsyncOS for SMTP Authentication

AsyncOS 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 appliance to make a secure connection to a relay server in configurations where the appliance is not at the edge of the network.

AsyncOS supports two methods to authenticate user credentials:

- You can use an LDAP directory.
- You can use a different SMTP server (SMTP Auth forwarding and SMTP Auth outgoing).

Figure 25-11 SMTP Auth Support: LDAP Directory Store or SMTP Server



Configured SMTP Authentication methods are then used to create SMTP Auth profiles via the `smtpauthconfig` command for use within HAT mail flow policies (see [Enabling SMTP Authentication on a Listener](#), page 25-36).

Related Topics

- [Configuring SMTP Authentication, page 25-33](#)
- [Configuring an SMTP Authentication Query, page 25-34](#)
- [SMTP Authentication via Second SMTP Server \(SMTP Auth with Forwarding\), page 25-35](#)
- [SMTP Authentication with LDAP, page 25-36](#)
- [Authenticating SMTP Sessions Using Client Certificates, page 25-39](#)
- [Outgoing SMTP Authentication, page 25-39](#)
- [Logging and SMTP Authentication, page 25-40](#)

Configuring SMTP Authentication

If you are going to authenticate with an LDAP server, select the SMTPAUTH query type on the Add or Edit LDAP Server Profile pages (or in the `ldapconfig` command) to create an SMTP Authentication query. For each LDAP server you configure, you can configure a SMTPAUTH query to be used as an SMTP Authentication profile.

There are two kinds of SMTP authentication queries: LDAP bind and Password as attribute. When you use password as attribute, the appliance will fetch the password field in the LDAP directory. The password may be stored in plain text, encrypted, or hashed. When you use LDAP bind, the appliance attempts to log into the LDAP server using the credentials supplied by the client.

Related Topics

- [Specifying a Password as Attribute, page 25-33](#)

Specifying a Password as Attribute

The convention in OpenLDAP, based on RFC 2307, is that the type of coding is prefixed in curly braces to the encoded password (for example, “{SHA}5en6G6MezRroT3XKqkdPOmY/BfQ=”). In this example, the password portion is a base64 encoding of a plain text password after application of SHA.

The appliance negotiates the SASL mechanism with the MUA before getting the password, and the appliance and the MUA decide on what method (LOGIN, PLAIN, MD5, SHA, SSHA, and CRYPT SASL mechanisms are supported). Then, the appliance queries the LDAP database to fetch a password. In LDAP, the password can have a prefix in braces.

- If there is no prefix, the appliance assumes that the password was stored in LDAP in plaintext.
- If there is a prefix, the appliance will fetch the hashed password, perform the hash on the username and/or password supplied by the MUA, and compare the hashed versions. The appliance supports SHA1 and MD5 hash types based on the RFC 2307 convention of prepending the hash mechanism type to the hashed password in the password field.
- Some LDAP servers, like the OpenWave LDAP server, do not prefix the encrypted password with the encryption type; instead, they store the encryption type as a separate LDAP attribute. In these cases, you can specify a default SMTP AUTH encryption method the appliance will assume when comparing the password with the password obtained in the SMTP conversation.

The appliance takes an arbitrary username from the SMTP Auth exchange and converts that to an LDAP query that fetches the clear or hashed password field. It will then perform any necessary hashing on the password supplied in the SMTP Auth credentials and compare the results with what it has retrieved from LDAP (with the hash type tag, if any, removed). A match means that the SMTP Auth conversation shall proceed. A failure to match will result in an error code.

Configuring an SMTP Authentication Query

Table 25-6 SMTP Auth LDAP Query Fields

Name	A name for the query.
Query String	You can select whether to authenticate via LDAP bind or by fetching the password as an attribute. Bind: Attempt to log into the LDAP server using the credentials supplied by the client (this is called an LDAP bind). Specify the maximum number of concurrent connections to be used by the SMTP Auth query. This number should not exceed the number specified in the LDAP server attributes above. Note, to avoid large number of session time-outs for bind authentication, increase the maximum number of concurrent connections here (typically nearly all of the connections can be assigned to SMTP Auth). A new connection is used for each bind authentication. The remainder of the connections are shared by the other LDAP query types. Password as Attribute: To authenticate by fetching passwords, specify the password in the SMTP Auth password attribute field below. Specify the LDAP query to use for either kind of authentication. Active Directory example query: <pre>(&(samaccountname={u})(objectCategory=person) (objectClass=user))</pre>
SMTP Auth Password Attribute	If you have selected “Authenticate by fetching the password as an attribute,” you can specify the password attribute here.

In the following example, the System Administration > LDAP page is used to edit the LDAP configuration named “PublicLDAP” to include an SMTPAUTH query. The query string (`uid={u}`) is constructed to match against `userPassword` attribute.

Figure 25-12 SMTP Authentication Query

☒ SMTP Authentication Query	
Name:	PublicLDAP.smtpauth
Query String:	{uid={u}}
	User Identity for Test Queries: Test Query
	Test SMTP Authentication Password: ?
Authentication Method:	☐ Authenticate via LDAP BIND Maximum number of concurrent connections for this query: 1 ☒ Authenticate by fetching the password as an attribute SMTP Authentication Password Attribute: userPassword

When an SMTPAUTH profile has been configured, you can specify that the listener uses that query for SMTP authentication.

SMTP Authentication via Second SMTP Server (SMTP Auth with Forwarding)

You can configure the appliance to verify the username and password that have been provided to another SMTP authenticated conversation with a different SMTP server.

The authenticating server is not the server that transfers mail; rather, it only responds to SMTP Authentication requests. When authentication has succeeded, the SMTP transfer of mail with the dedicated mail server can proceed. This feature is sometimes referred to as “SMTP Authentication with forwarding” because only the credentials are forwarded (or “proxied”) to another SMTP server for authentication.

Procedure

-
- Step 1** Choose Network > SMTP Authentication.
 - Step 2** Click **Add Profile...**
 - Step 3** Enter a unique name for the SMTP authentication profile.
 - Step 4** For the **Profile Type**, select **Forward**.
 - Step 5** Click **Next**.
 - Step 6** Enter the hostname/IP address and port of the forwarding server. Select a forwarding interface to use for forwarding authentication requests. Specify the number of maximum simultaneous connections. Then, you can configure whether TLS is required for connections from the appliance to the forwarding server. You can also select a SASL method to use (PLAIN or LOGIN), if available. This selection is configured for each forwarding server.
 - Step 7** Submit and commit your changes.
 - Step 8** After creating the authentication profile, you can enable the profile on a listener. See [Enabling SMTP Authentication on a Listener, page 25-36](#) for more information.
-

SMTP Authentication with LDAP

To create an LDAP-based SMTP Authentication profile, you must have previously created an SMTP Authentication query in conjunction with an LDAP server profile using the System Administration > LDAP page. You can then use this profile to create an SMTP Authentication profile. For more information about creating an LDAP profile, see [Understanding LDAP Queries, page 25-2](#).

Procedure

-
- Step 1** Choose Network > SMTP Authentication.
 - Step 2** Click **Add Profile**.
 - Step 3** Enter a unique name for the SMTP authentication profile.
 - Step 4** For the Profile Type, select **LDAP**.
 - Step 5** Click **Next**.
 - Step 6** Select the LDAP query you would like to use for this authentication profile.
 - Step 7** Select a default encryption method from the drop-down menu. You can select from SHA, Salted SHA, Crypt, Plain, or MD5. If your LDAP servers prefix an encrypted password with the encryption type, leave 'None' selected. If your LDAP server saves the encryption type as a separate entity (OpenWave LDAP servers, for example), then select an encryption method from the menu. The default encryption setting will not be used if the LDAP query is using bind.
 - Step 8** Click **Finish**.
 - Step 9** Submit and commit your changes.
 - Step 10** After creating the authentication profile, you can enable the profile on a listener. See [Enabling SMTP Authentication on a Listener, page 25-36](#) for more information.
-

Related Topics

- [Enabling SMTP Authentication on a Listener, page 25-36](#)

Enabling SMTP Authentication on a Listener

After using the Network > SMTP Authentication page to create an SMTP authentication “profile” that specifies the type of SMTP authentication you want to perform (LDAP-based or SMTP forwarding-based), you must associate that profile with a listener using the Network > Listeners page (or the `listenerconfig` command).



Note

An authenticated user is granted RELAY connection behavior within their current Mail Flow Policy.



Note

You may specify more than one forwarding server in a profile. SASL mechanisms CRAM-MD5 and DIGEST-MD5 are not supported between the appliance and a forwarding server.

In the following example, the listener “InboundMail” is edited to use the SMTPAUTH profile configured via the Edit Listener page:

Figure 25-13 Selecting an SMTP Authentication Profile via the Edit Listener page

Listener Settings	
Name:	IncomingMail
Type of Listener:	Public
Interface:	Data 1 TCP Port: 25
Bounce Profile:	Default
Disclaimer Above:	None Disclaimer text will be applied above the message body.
Disclaimer Below:	None Disclaimer text will be applied below the message body.
SMTP Authentication Profile:	forwarding_based
Certificate:	test
▶ SMTP Address Parsing Options:	Optional settings for controlling parsing in SMTP "MAIL FROM" and "RCPT TO" commands.
▶ Advanced:	Optional settings for customizing the behavior of the Listener

Once a listener is configured to use the profile, the Host Access Table default settings can be changed so that the listener allows, disallows, or requires SMTP Authentication:

Figure 25-14 Enabling SMTP Authentication on a Mail Flow Policy

Encryption and Authentication:	TLS:	☐ Use Default (Off) ☒ Off ☐ Preferred ☐ Required
	① SMTP Authentication:	☒ Use Default (Off) ☐ Off ☐ Preferred ☐ Required
	② If Both TLS and SMTP Authentication are enabled:	☐ Require TLS To Offer SMTP Authentication

Number	Description
1.	The SMTP Authentication field provides listener-level control for SMTP authentication. If you select “No,” authentication will not be enabled on the listener, regardless of any other SMTP authentication settings you configure.
2.	If “Required” is selected in the second prompt (SMTP Authentication:), no AUTH keyword will be issued until TLS is negotiated (after the client issues a second EHLO command).

Related Topics

- [SMTP Authentication and HAT Policy Settings, page 25-37](#)
- [HAT Delayed Rejection, page 25-38](#)

SMTP Authentication and HAT Policy Settings

Because senders are grouped into the appropriate sender group before the SMTP Authentication negotiation begins, Host Access Table (HAT) settings, are not affected. When a remote mail host connects, the appliance first determines which sender group applies and imposes the Mail Policy for that sender group. For example, if a remote MTA “suspicious.com” is in your SUSPECTLIST sender group, the THROTTLE policy will be applied, regardless of the results of “suspicious.com’s” SMTPAUTH negotiation.

However, senders that do authenticate using SMTPAUTH are treated differently than “normal” senders. The connection behavior for successful SMTPAUTH sessions changes to “RELAY,” effectively bypassing the Recipient Access Table (RAT) and LDAPACCEPT. This allows the sender to relay messages through the appliance. As stated, any Rate Limiting or throttling that applies will remain in effect.

HAT Delayed Rejection

When HAT Delayed Rejection is configured, connections that would get dropped based on the HAT Sender Group and Mail Flow Policy configuration can still authenticate successfully and get the RELAY mail flow policy granted.

You can configure delayed rejection using the `listenerconfig --> setup` CLI command. This behavior is disabled by default.

The following table shows how to configure delayed rejection for HAT.

```
example.com> listenerconfig
```

```
Currently configured listeners:
```

1. listener1 (on main, 172.22.138.17) QMQP TCP Port 628 Private
2. listener2 (on main, 172.22.138.17) SMTP TCP Port 25 Private

```
Choose the operation you want to perform:
```

- NEW - Create a new listener.
- EDIT - Modify a listener.
- DELETE - Remove a listener.
- SETUP - Change global settings.

```
[> setup
```

```
Enter the global limit for concurrent connections to be allowed across all listeners.
```

```
[300]>
```

```
[...]
```

```
By default HAT rejected connections will be closed with a banner
```

```
message at the start of the SMTP conversation. Would you like to do the rejection at the  
message recipient level instead for more detailed logging of rejected mail?
```

```
[N]> y
```

Do you want to modify the SMTP RCPT TO reject response in this case?

[N]> **y**

Enter the SMTP code to use in the response. 550 is the standard code.

[550]> **551**

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

Sender rejected due to local mail policy.

Contact your mail admin for assistance.

Authenticating SMTP Sessions Using Client Certificates

The Email Security appliance supports the use of client certificates to authenticate SMTP sessions between the Email Security appliance and users' mail clients.

When creating an SMTP authentication profile, you select the Certificate Authentication LDAP query to use for verifying the certificate. You can also specify whether the Email Security appliance falls back to the SMTP AUTH command to authenticate the user if a client certificate isn't available.

If your organization uses client certificates to authenticate users, you have the option of using the SMTP Authentication query to check whether a user who doesn't have a client certificate can send mail as long as their record specifies that it's allowed.

See [Authenticating SMTP Sessions Using Client Certificates](#) for more information.

Outgoing SMTP Authentication

SMTP Authentication can also be used to provide validation for an outbound mail relay, using a username and password. Create an 'outgoing' SMTP authentication profile and then attach the profile to an SMTP route for the ALL domain. On each mail delivery attempt, the appliance will log on to the upstream mail relay with the necessary credentials. SMTP authentication supports the following authorization protocols: PLAIN and LOGIN.

Procedure

-
- Step 1** Choose **Network > SMTP Authentication**.
 - Step 2** Click **Add Profile**.
 - Step 3** Enter a unique name for the SMTP authentication profile.
 - Step 4** For the Profile Type, select **Outgoing**.
 - Step 5** Click **Next**.

- Step 6** Enter an authentication username and password for the authentication profile.
 - Step 7** Click **Finish**.
 - Step 8** Choose **Network > SMTP Routes**.
 - Step 9** Click the **All Other Domains** link in the **Receiving Domain** column of the table.
 - Step 10** Enter the name of the Destination Host for the SMTP route. This is the hostname of your external mail relay used to deliver outgoing mail.
 - Step 11** Select the outgoing SMTP authentication profile from the drop-down menu.
 - Step 12** Submit and commit your changes.
-

Logging and SMTP Authentication

The following events will be logged in the mail logs when the SMTP Authentication mechanism (either LDAP-based, SMTP forwarding server based, or SMTP outgoing) is configured on the appliance:

- [Informational] Successful SMTP Authentication attempts — including the user authenticated and the mechanism used. (No plaintext passwords will be logged.)
- [Informational] Unsuccessful SMTP Authentication attempts — including the user authenticated and the mechanism used.
- [Warning] Inability to connect to the authentication server — including the server name and the mechanism.
- [Warning] A time-out event when the forwarding server (talking to an upstream, injecting appliance) times out while waiting for an authentication request.

Configuring External LDAP Authentication for Users

You can configure the appliance to use an LDAP directory on your network to authenticate users by allowing them to log in with their LDAP usernames and passwords. After you configure the authentication queries for the LDAP server, enable the appliance to use external authentication on the System Administration > Users page in the GUI (or use the `userconfig` command in the CLI).

Procedure

- Step 1** **Create a query to find user accounts.** In an LDAP server profile, create a query to search for user accounts in the LDAP directory.
- Step 2** **Create group membership queries.** Create a query to determine if a user is a member of a directory group.
- Step 3** **Set up external authentication to use the LDAP server.** Enable the appliance to use the LDAP server for user authentication and assign user roles to the groups in the LDAP directory. For more information, see “Adding Users” in the “Distributing Administrative Tasks” chapter.

**Note**

Use the Test Query button on the LDAP page (or the `ldaptest` command) to verify that your queries return the expected results. For more information, see [Testing LDAP Queries, page 25-17](#).

Related Topics

- [User Accounts Query, page 25-41](#)
- [Group Membership Queries, page 25-41](#)

User Accounts Query

To authenticate external users, AsyncOS uses a query to search for the user record in the LDAP directory and the attribute that contains the user's full name. Depending on the server type you select, AsyncOS enters a default query and a default attribute. You can choose to have your appliance deny users with expired accounts if you have attributes defined in RFC 2307 in your LDAP user records (`shadowLastChange`, `shadowMax`, and `shadowExpire`). The base DN is required for the domain level where user records reside.

[Table 25-7](#) shows the default query string and full username attribute that AsyncOS uses when it searches for a user account on an Active Directory server.

Table 25-7 **Default User Account Query String and Attribute: Active Directory**

Server Type	Active Directory
Base DN	[blank] (You need to use a specific base DN to find the user records.)
Query String	<code>(&(objectClass=user)(sAMAccountName={u}))</code>
Attribute containing the user's full name	<code>displayName</code>

[Table 25-8](#) shows the default query string and full username attribute that AsyncOS uses when it searches for a user account on an OpenLDAP server.

Table 25-8 **Default User Account Query String and Attribute: OpenLDAP**

Server Type	OpenLDAP
Base DN	[blank] (You need to use a specific base DN to find the user records.)
Query String	<code>(&(objectClass=posixAccount)(uid={u}))</code>
Attribute containing the user's full name	<code>gecos</code>

Group Membership Queries

AsyncOS also uses a query to determine if a user is a member of a directory group. Membership in a directory group membership determines the user's permissions within the system. When you enable external authentication on the System Administration > Users page in the GUI (or `userconfig` in the

CLI), you assign user roles to the groups in your LDAP directory. User roles determine the permissions that users have in the system, and for externally authenticated users, the roles are assigned to directory groups instead of individual users. For example, you can assign users in the IT directory group the Administrator role and users in the Support directory group to the Help Desk User role.

If a user belongs to multiple LDAP groups with different user roles, AsyncOS grants the user the permissions for the most restrictive role. For example, if a user belongs to a group with Operator permissions and a group with Help Desk User permissions, AsyncOS grants the user the permissions for the Help Desk User role.

When you configure the LDAP profile to query for group membership, enter the base DN for the directory level where group records can be found, the attribute that holds the group member's username, and the attribute that contains the group name. Based on the server type that you select for your LDAP server profile, AsyncOS enters default values for the username and group name attributes, as well default query strings.

**Note**

For Active Directory servers, the default query string to determine if a user is a member of a group is `(&(objectClass=group)(member={u}))`. However, if your LDAP schema uses distinguished names in the “memberof” list instead of usernames, you can use `{dn}` instead of `{u}`.

Table 25-9 shows the default query strings and attributes that AsyncOS uses when it searches for group membership information on an Active Directory server.

Table 25-9 Default Group Membership Query Strings and Attribute: Active Directory

Server Type	Active Directory
Base DN	[blank] (You need to use a specific base DN to find the group records.)
Query string to determine if a user is a member of a group	<code>(&(objectClass=group)(member={u}))</code> Note If your LDAP schema uses distinguished names in the <code>memberof</code> list instead of usernames, you can replace <code>{u}</code> with <code>{dn}</code> .
Attribute that holds each member's username (or a DN for the user's record)	<code>member</code>
Attribute that contains the group name	<code>cn</code>

Table 25-10 shows the default query strings and attributes that AsyncOS uses when it searches for group membership information on an OpenLDAP server.

Table 25-10 Default Group Membership Query Strings and Attributes: OpenLDAP

Server Type	OpenLDAP
Base DN	[blank] (You need to use a specific base DN to find the group records.)
Query string to determine if a user is a member of a group	<code>(&(objectClass=posixGroup)(memberUid={u}))</code>
Attribute that holds each member's username (or a DN for the user's record)	<code>memberUid</code>
Attribute that contains the group name	<code>cn</code>

Authenticating End-Users of the Spam Quarantine

Spam quarantine end-user authentication queries validate users when they log in to the Spam Quarantine. The token {u} specifies the user (it represents the user's login name). The token {a} specifies the user's email address. The LDAP query does not strip "SMTP:" from the email address; AsyncOS strips that portion of the address.

If you want the Spam Quarantine to use an LDAP query for end-user access, check the "Designate as the active query" check box. If there is an existing active query, it is disabled. When you open the System Administration > LDAP page, an asterisk (*) is displayed next to the active queries.

Based on the server type, AsyncOS uses one of the following default query strings for the end-user authentication query:

- **Active Directory:** (sAMAccountName={u})
- **OpenLDAP:** (uid={u})
- **Unknown or Other:** [Blank]

By default, the primary email attribute is proxyAddresses for Active Directory servers and mail for OpenLDAP servers. You can enter your own query and email attributes. To create the query from the CLI, use the isgauth subcommand of the ldapconfig command.



Note

If you want users to log in with their full email address, use (mail=smtp:{a}) for the Query String.

Related Topics

- [Sample Active Directory End-User Authentication Settings, page 25-43](#)
- [Sample OpenLDAP End-User Authentication Settings, page 25-44](#)
- [Configuring End-User Access to the Spam Quarantine, page 31-17](#)
- [Configuring End-User Access to the Spam Quarantine, page 31-17](#)

Sample Active Directory End-User Authentication Settings

This section shows sample settings for an Active Directory server and the end-user authentication query. This example uses password authentication for the Active Directory server, the mail and proxyAddresses email attributes, and the default query string for end-user authentication for Active Directory servers.

Table 25-11 Example LDAP Server and Spam Quarantine End-User Authentication Settings: Active Directory

Authentication Method	Use Password (Need to create a low-privilege user to bind for searching, or configure anonymous searching.)
Server Type	Active Directory
Port	3268
Base DN	[Blank]
Connection Protocol	[Blank]
Query String	(sAMAccountName={u})
Email Attribute(s)	mail,proxyAddresses

Sample OpenLDAP End-User Authentication Settings

This section shows sample settings for an OpenLDAP server and the end-user authentication query. This example uses anonymous authentication for the OpenLDAP server, the `mail` and `mailLocalAddress` email attributes, and the default query string for end-user authentication for OpenLDAP servers.

Table 25-12 Example LDAP Server and Spam Quarantine End-User Authentication Settings: OpenLDAP

Authentication Method	Anonymous
Server Type	OpenLDAP
Port	389
Base DN	[Blank] (Some older schemas will want to use a specific Base DN.)
Connection Protocol	[Blank]
Query String	(uid={u})
Email Attribute(s)	mail,mailLocalAddress

Spam Quarantine Alias Consolidation Queries

If you use spam notifications, the spam quarantine alias consolidation query consolidates the email aliases so that recipients do not receive quarantine notices for each alias. For example, a recipient might receive mail for the following email addresses: `john@example.com`, `jsmith@example.com`, and `john.smith@example.com`. When you use alias consolidation, the recipient receives a single spam notification at a chosen primary email address for messages sent to all of the user's aliases.

To consolidate messages to a primary email address, create a query to search for a recipient's alternate email aliases, and then enter the attribute for the recipient's primary email address in the Email Attribute field.

If you want the Spam Quarantine to use an LDAP query for spam notifications, check the "Designate as the active query" check box. If there is an existing active query, it is disabled. When you open the System Administration > LDAP page, an asterisk (*) is displayed next to the active queries.

For Active Directory servers, the default query string is `(|(proxyAddresses={a})(proxyAddresses=smtp:{a}))` and the default email attribute is `mail`. For OpenLDAP servers, the default query string is `(mail={a})` and the default email attribute is `mail`. You can define your own query and email attributes, including multiple attributes separated by commas. If you enter more than one email attribute, Cisco recommends entering a unique attribute that uses a single value, such as `mail`, as the first email attribute instead of an attribute with multiple values that can change, such as `proxyAddresses`.

To create the query in the CLI, use the `isqalias` subcommand of the `ldapconfig` command.

Related Topics

- [Sample Active Directory Alias Consolidation Settings, page 25-45](#)
- [Sample OpenLDAP Alias Consolidation Settings, page 25-45](#)

Sample Active Directory Alias Consolidation Settings

This section shows sample settings for an Active Directory server and the alias consolidation query. This example uses anonymous authentication for the Active Directory server, a query string for alias consolidation for Active Directory servers, and the `mail` email attribute.

Table 25-13 Example LDAP Server and Spam Quarantine Alias Consolidation Settings: Active Directory

Authentication Method	Anonymous
Server Type	Active Directory
Port	3268
Base DN	[Blank]
Connection Protocol	Use SSL
Query String	((mail={a}) (mail=smtpp:{a}))
Email Attribute	mail

Sample OpenLDAP Alias Consolidation Settings

This section shows sample settings for an OpenLDAP server and the alias consolidation query. This example uses anonymous authentication for the OpenLDAP server, a query string for alias consolidation for OpenLDAP servers, and the `mail` email attribute.

Table 25-14 Example LDAP Server and Spam Quarantine Alias Consolidation Settings: OpenLDAP

Authentication Method	Anonymous
Server Type	OpenLDAP
Port	389
Base DN	[Blank] (Some older schemas will want to use a specific Base DN.)
Connection Protocol	Use SSL
Query String	(mail={a})
Email Attribute	mail

Identifying a Sender's User Distinguished Name for RSA Enterprise Manager

the Email Security appliance must include the complete distinguished names for the message senders when it sends DLP incident data to Enterprise Manager. To acquire the sender name for Enterprise Manager, create a user distinguished name query for your LDAP server and add the query to the listeners that send outgoing messages on your Email Security appliance. The Email Security appliance only uses this query when RSA Enterprise Manager is enabled for DLP. Otherwise, it does not appear as an option for the server profile.

Related Topics

- [Sample User Distinguished Name Settings, page 25-46](#)
- [Configuring AsyncOS To Work With Multiple LDAP Servers, page 25-46](#)
- [Testing Servers and Queries, page 25-47](#)
- [Failover, page 25-47](#)
- [Load Balancing, page 25-48](#)

Sample User Distinguished Name Settings

This section shows sample settings for an Active Directory server and the user distinguished name query. This example uses anonymous authentication for the Active Directory server and a query string for user distinguished name retrieval for Active Directory servers.

Table 25-15 *Example LDAP Server and Spam Quarantine Alias Consolidation Settings: Active Directory*

Authentication Method	Anonymous
Server Type	Active Directory
Port	3268
Base DN	[Blank]
Connection Protocol	Use SSL
Query String	(proxyAddresses=smtp:{a})

Configuring AsyncOS To Work With Multiple LDAP Servers

When you configure an LDAP profile, you can configure the appliance to connect to a list of multiple LDAP servers. To use multiple LDAP servers, you must configure LDAP servers to contain the same information, use the same structure, and use the same authentication information. (third party products exist that can consolidate the records).

When you configure the appliance to connect to redundant LDAP servers, you can configure the LDAP configuration for failover or load balancing.

You can use multiple LDAP servers to achieve the following results:

- **Failover.** When you configure the LDAP profile for failover, the appliance fails over to the next LDAP server in the list if it cannot connect to the first LDAP server.
- **Load Balancing.** When you configure the LDAP profile for load balancing, the appliance distributes connections across the list of LDAP servers when it performs LDAP queries.

You can configure redundant LDAP servers from the System Administration > LDAP page or from the CLI `ldapconfig` command.

Testing Servers and Queries

Use the **Test Server(s)** button on the Add (or Edit) LDAP Server Profile page (or the `test` subcommand in the CLI) to test the connection to an LDAP server. If you use multiple LDAP servers, AsyncOS tests each server and displays individual results for each server. AsyncOS will also test the query on each LDAP server and display the individual results.

Failover

To ensure that LDAP queries are resolved, you can configure your LDAP profile for failover.

The appliance attempts to connect to the first server in the list of LDAP servers for a specified period of time. If the appliance cannot connect to the first LDAP server in the list, the appliance attempts to connect to the next LDAP server in the list. By default, the appliance always attempts to connect to the first server in the list, and it attempts to connect to each subsequent server in the order they are listed. To ensure that the appliance connects to your primary LDAP server by default, ensure that you enter it as the first server in your list of LDAP servers.

If the appliance connects to a second or subsequent LDAP server, it remains connected to that server until it reaches a timeout period. After it reaches the timeout, it attempts to reconnect to the first server in the list.

Related Topics

- [Configuring the Appliance for LDAP Failover, page 25-47](#)

Configuring the Appliance for LDAP Failover

To configure the appliance for LDAP failover, complete the following steps in the GUI:

Procedure

Step 1 From System Administration > LDAP, select the LDAP server profile you want to edit.

Step 2 From the LDAP server profile, configure the following settings:

The screenshot shows the 'LDAP Server Settings' window with the 'Server Attributes' tab selected. The 'LDAP Server Configuration Name' is 'example.com'. The 'Host Name(s)' field contains 'ldapserver1.example.com, ldapserver2.example.com, ldapserver3.example.com' with a note 'Separate multiple entries with commas.' The 'Maximum number of simultaneous connections for all hosts' is set to '10'. Under 'Multiple host options', the 'Failover connections in the order listed' radio button is selected.

Number	Description
1	List LDAP Servers.
2	Configure Maximum Connections.
3	Select Failover Mode.

Step 3 Configure other LDAP settings and commit the changes.

Load Balancing

To distribute LDAP connections among a group of LDAP servers, you can configure your LDAP profile for load balancing.

When you configure your LDAP profile for load balancing, the appliance distributes connections among the LDAP servers listed. If a connection fails or times out, the appliance determines which LDAP servers are available and reconnects to available servers. The appliance determines the number of simultaneous connections to establish based on the maximum number of connections you configure.

If one of the listed LDAP servers does not respond, the appliance distributes the connection load among the remaining LDAP servers.

Reliance Topics

- [Configuring the Appliance for Load Balancing, page 25-48](#)

Configuring the Appliance for Load Balancing

Procedure

Step 1 From System Administration > LDAP, select the LDAP server profile you want to edit.

Step 2 From the LDAP server profile, configure the following settings:

Server Attributes	
LDAP Server Configuration Name:	example.com
① Host Name(s):	ldapservice1.example.com, ldapservice2.example.com, ldapservice3.example.com Separate multiple entries with commas.
	Maximum number of simultaneous connections for all hosts: 10 ②
③	Multiple host options: ☒ Load-balance connections among all hosts listed ☐ Failover connections in the order listed

Number	Description
1	List LDAP Servers
2	Configure Maximum Connections
3	Select Load Balancing Mode

Step 3 Configure other LDAP settings and commit the changes.



Authenticating SMTP Sessions Using Client Certificates

- [Overview of Certificates and SMTP Authentication, page 26-49](#)
- [Checking the Validity of a Client Certificate, page 26-51](#)
- [Authenticating a User Using an LDAP Directory, page 26-52](#)
- [Authenticating an SMTP Connection Over TLS Using a Client Certificate, page 26-52](#)
- [Establishing a TLS Connection from the Appliance, page 26-53](#)
- [Updating a List of Revoked Certificates, page 26-54](#)

Overview of Certificates and SMTP Authentication

The Email Security appliance supports the use of client certificates to authenticate SMTP sessions between the Email Security appliance and users' mail clients. The Email Security appliance can request a client certificate from a user's mail client when the application attempts to connect to the appliance to send messages. When the appliance receives the client certificate, it verifies that the certificate is valid, has not expired, and has not been revoked. If the certificate is valid, the Email Security appliance allows an SMTP connection from the mail application over TLS.

Organizations that require their users to use a Common Access Card (CAC) for their mail clients can use this feature to configure the Email Security appliance to request a certificate that the CAC and ActivClient middleware application will provide to the appliance.

You can configure the Email Security appliance to require users to provide a certificate when sending mail, but still allow exceptions for certain users. For these users, you can configure the appliance to use the SMTP authentication LDAP query to authenticate the user.

Users must configure their mail client to send messages through a secure connection (TLS) and accept a server certificate from the appliance.

Related Topics

- [How to Authenticate a User with a Client Certificate, page 26-50](#)
- [How to Authenticate a User with an SMTP Authentication LDAP Query, page 26-50](#)
- [How to Authenticate a User with an LDAP SMTP Authentication Query if the Client Certificate is Invalid, page 26-50](#)

How to Authenticate a User with a Client Certificate

Table 26-1 *How to Authenticate a User with a Client Certificate*

	Do This	More Info
Step 1	Define a certificate query for your LDAP server.	Checking the Validity of a Client Certificate, page 26-51
Step 2	Create a certificate-based SMTP authentication profile.	Authenticating an SMTP Connection Over TLS Using a Client Certificate, page 26-52
Step 3	Configure a listener to use the certificate SMTP authentication profile.	Listening for Connection Requests by Creating a Listener via the GUI, page 5-8
Step 4	Modify the RELAYED mail flow policy to require TLS, a client certificate, and SMTP authentication.	Establishing a TLS Connection from the Appliance, page 26-53

How to Authenticate a User with an SMTP Authentication LDAP Query

Table 26-2 *How to Authenticate a User with an SMTP Authenticate LDAP Query*

	Do This	More Info
Step 1	Define an SMTP authentication query for your server that uses an allowance query string and Bind for the authentication method.	Authenticating a User Using an LDAP Directory, page 26-52
Step 2	Create an LDAP-based SMTP authentication profile.	Configuring AsyncOS for SMTP Authentication, page 25-32
Step 3	Configure a listener to use the LDAP SMTP authentication profile.	If the user is not allowed to use LDAP-based SMTP authentication for their connection, you can select whether the appliance rejects the connection or temporarily allows it while logging all activity.
Step 4	Modify the RELAYED mail flow policy to require TLS and SMTP authentication.	Establishing a TLS Connection from the Appliance, page 26-53

How to Authenticate a User with an LDAP SMTP Authentication Query if the Client Certificate is Invalid

Table 26-3 *How to Authenticate a User with a Client Certificate or an LDAP SMTP Authentication Query*

	Do This	More Info
Step 1	Define an SMTP authentication query for your server that uses an allowance query string and Bind for the authentication method.	Authenticating a User Using an LDAP Directory, page 26-52
Step 2	Define a certificate-based query for your LDAP server.	Checking the Validity of a Client Certificate, page 26-51

Table 26-3 *How to Authenticate a User with a Client Certificate or an LDAP SMTP Authentication Query*

	Do This	More Info
Step 3	Create a certificate-based SMTP authentication profile	Authenticating an SMTP Connection Over TLS Using a Client Certificate, page 26-52
Step 4	Create an LDAP SMTP authentication profile.	Configuring AsyncOS for SMTP Authentication, page 25-32
Step 5	Configure a listener to use the certificate SMTP authentication profile.	Listening for Connection Requests by Creating a Listener via the GUI, page 5-8
Step 6	1. Modify the RELAYED mail flow policy to use the following settings: • TLS Preferred • SMTP authentication required • Require TLS for SMTP authentication	Establishing a TLS Connection from the Appliance, page 26-53

Checking the Validity of a Client Certificate

The Certificate Authentication LDAP query checks the validity of a client certificate in order to authenticate an SMTP session between the user's mail client and the Email Security appliance. When creating this query, you select a list of certificate fields for authentication, specify the User ID attribute (the default is `uid`), and enter the query string.

For example, a query string that searches for the certificate's common name and serial number may look like `(&(objectClass=posixAccount)(cacn={cn})(cacserial={sn}))`. After you have created the query, you can use it in a Certificate SMTP Authentication Profile. This LDAP query supports OpenLDAP, Active Directory, and Oracle Directory.

See [Chapter 25, "LDAP Queries"](#) for more information on configuring LDAP servers.

Procedure

-
- Step 1** Select **System Administration > LDAP**.
 - Step 2** Create a new LDAP profile. See [Creating LDAP Server Profiles to Store Information About the LDAP Server, page 25-5](#) for more information.
 - Step 3** Check the **Certificate Authentication Query** checkbox.
 - Step 4** Enter the query name.
 - Step 5** Enter the query string to authenticate the user's certificate. For example, `(&(objectClass=user)(cn={cn}))`.
 - Step 6** Enter the user ID attribute, such as `sAMAccountName`.
 - Step 7** Submit and commit your changes.
-

Authenticating a User Using an LDAP Directory

The SMTP Authentication LDAP query has an Allowance Query String that allows the Email Security appliance to check whether the user's mail client is allowed to send mail through the appliance based on the user's record in the LDAP directory. This allows users who don't have a client certificate to send mail as long as their record specifies that it's allowed.

You can also filter out results based on other attributes. For example, the query string `(&(uid={u})(|(! (caccn=*)) (cacexempt=*) (cacemergency>={t})))` checks to see if any of the following conditions are true for the user:

- CAC is not issued to the user (`caccn=*`)
- CAC is exempt (`cacexempt=*`)
- the time period that a user may temporarily send mail without a CAC expires in the future (`cacemergency>={t}`)

See [Configuring AsyncOS for SMTP Authentication, page 25-32](#) for more information on using the SMTP Authentication query.

Procedure

-
Step 1	Select System Administration > LDAP .
Step 2	Define an LDAP profile. See Creating LDAP Server Profiles to Store Information About the LDAP Server, page 25-5 for more information.
Step 3	Define an SMTP authentication query for the LDAP profile.
Step 4	Check the SMTP Authentication Query checkbox.
Step 5	Enter the query name.
Step 6	Enter the string to query for the user's ID. For example, <code>(uid={u})</code> .
Step 7	Select LDAP BIND for the authentication method.
Step 8	Enter an allowance query string. For example, <code>(&(uid={u})((! (caccn=*)) (cacexempt=*) (cacemergency>={t})))</code> .
Step 9	Submit and commit your changes.
-

Authenticating an SMTP Connection Over TLS Using a Client Certificate

The certificate-based SMTP authentication profile allows the Email Security appliance to authenticate an SMTP connection over TLS using a client certificate. When creating the profile, you select the Certificate Authentication LDAP query to use for verifying the certificate. You can also specify whether the Email Security appliance falls back to the SMTP AUTH command to authenticate the user if a client certificate isn't available.

For information on authenticating an SMTP connection by using LDAP, see [Configuring AsyncOS for SMTP Authentication, page 25-32](#).

Procedure

-
- Step 1** Select **Network > SMTP Authentication**.
- Step 2** Click **Add Profile**.
- Step 3** Enter the name for the SMTP authentication profile.
- Step 4** Select **Certificate** for the Profile Type.
- Step 5** Click **Next**.
- Step 6** Enter the profile name.
- Step 7** Select the certificate LDAP query you want to use with this SMTP authentication profile.



Note Do not select the option to allow the SMTP AUTH command if a client certificate is not available.

- Step 8** Click **Finish**.
- Step 9** Submit and commit your changes.
-

Establishing a TLS Connection from the Appliance

The Verify Client Certificate option in the RELAYED mail flow policy directs the Email Security appliance to establish a TLS connection to the user's mail application if the client certificate is valid. If you select this option for the TLS Preferred setting, the appliance still allows a non-TLS connection if the user doesn't have a certificate, but rejects a connection if the user has an invalid certificate. For the TLS Required setting, selecting this option requires the user to have a valid certificate in order for the appliance to allow the connection.

To authenticate a user's SMTP session with a client certificate, select the following settings:

- TLS - Required
- Verify Client Certificate
- Require SMTP Authentication



Note Although SMTP authentication is required, the Email Security appliance will not use the SMTP authentication LDAP query because it is using certificate authentication.

To authenticate a user's SMTP session using the SMTP authentication query instead of a client certificate, select the following settings for the RELAYED mail flow policy:

- TLS - Required
- Require SMTP Authentication

If you require the Email Security appliance to ask for a client certificate from certain users while allowing LDAP-based SMTP authentication from others, select the following settings for the RELAYED mail flow policy:

- TLS - Preferred

- Require SMTP Authentication
- Require TLS to Offer SMTP Authentication

Updating a List of Revoked Certificates

The Email Security appliance checks a list of revoked certificates (called a Certificate Revocation List) as part of its certificate verification to make sure that the user's certificate hasn't been revoked. You keep an up-to-date version of this list on a server and the Email Security appliance downloads it on a schedule that you create.

Procedure

-
- Step 1** Go to **Network > CRL Sources**.
 - Step 2** Enable CRL checking for SMTP TLS connections:
 - a. Click **Edit Settings** under **Global Settings**.
 - b. Select the checkbox for **CRL check for inbound SMTP TLS**.
 - c. (Optional) Select the checkbox for **CRL check for inbound SMTP TLS**.
 - d. Submit your change.
 - Step 3** Click **Add CRL Source**.
 - Step 4** Enter a name for the CRL source.
 - Step 5** Select the file type. This can be either **ASN.1** or **PEM**.
 - Step 6** Enter the URL for the primary source for the file, including the filename. For example, `https://crl.example.com/certs.crl`
 - Step 7** Optionally, enter the URL for a secondary source in case the appliance cannot contact the primary source.
 - Step 8** Specify a schedule for downloading the CRL source.
 - Step 9** Enable the CRL source.
 - Step 10** Submit and commit your changes.
-



FIPS Management

- [FIPS Management Overview, page 27-1](#)
- [Configuration Changes in FIPS Mode, page 27-1](#)
- [Switching the Appliance to FIPS Mode, page 27-2](#)
- [Encrypting Sensitive Data in FIPS Mode, page 27-3](#)
- [Checking FIPS Mode Compliance, page 27-4](#)
- [Managing Certificates and Keys, page 27-4](#)
- [Managing Keys for DKIM Signing and Verification, page 27-5](#)

FIPS Management Overview

The Federal Information Processing Standard (FIPS) 140 is a publicly announced standard developed jointly by the United States and Canadian federal governments specifying requirements for cryptographic modules that are used by government agencies to protect sensitive but unclassified information. The Cisco IronPort Email Security appliance uses the Cisco SSL Cryptographic Toolkit to achieve FIPS 140-2 Level 1 compliance.

The Cisco SSL Cryptographic Toolkit is a GGSG-approved cryptography suite that includes Cisco SSL, which is an enhanced version of OpenSSL's FIPS support, and the FIPS-compliant Cisco Common Cryptography Module. The Cisco Common Cryptography Module is a software library that Email Security appliance uses for FIPS-validated cryptographic algorithms for protocols such as SSH.

Configuration Changes in FIPS Mode

The Email Security appliance uses Cisco SSL and FIPS-compliant certificates for communication when the appliance is in FIPS mode. See [Switching the Appliance to FIPS Mode, page 27-2](#) for more information.



Note

As part of FIPS compliance, AsyncOS for Email does not support SSH version 1.

To be FIPS Level 1 compliant, the Email Security appliance makes the following changes to your configuration:

- **SMTP receiving and delivery.** Incoming and outgoing SMTP conversations over TLS between a public listener on the Email Security appliance and a remote host use TLS version 1 and FIPS cipher suites. You can modify the cipher suites using `sslconfig` when in FIPS mode. TLS v1 is the only version of TLS supported in FIPS mode.
- **Web interface.** HTTPS sessions to the Email Security appliance's web interface use TLS version 1 and FIPS cipher suites. This also includes HTTPS sessions to the IronPort Spam Quarantine and other IP interfaces. You can modify the cipher suites using `sslconfig` when in FIPS mode.
- **Certificates.** FIPS mode restricts the kinds of certificates used by the appliances. Certificates must use one of the following signature algorithms: SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512 and RSA keys of the size 2048 bits. The appliance will not import certificates that do not use one of these algorithms. The appliance cannot be switched to FIPS mode if it has any non-compliant certificates in use. It will display an error message instead. See [Managing Certificates and Keys, page 27-4](#) for more information.
- **DKIM signing and verification.** RSA keys used for DKIM signatures and verification must be 2048 bits in length. The appliance cannot be switched to FIPS mode if it has any non-compliant RSA keys in use. It will display an error message instead. When verifying a DKIM signature, the appliance returns a permanent failure if the signature does not use a FIPS-compliant key. See [Chapter 27, "Managing Keys for DKIM Signing and Verification"](#)
- **LDAPs.** TLS transactions between the Email Security appliance and LDAP servers, including using an LDAP server for external authentication, use TLS version 1 and FIPS cipher suites. If the LDAP server uses MD5 hashes to store passwords, the SMTP authentication query will fail because MD5 is not FIPS-compliant.
- **Logs.** SSH2 is the only allowed protocol for pushing logs via SCP. For error messages related to FIPS management, read the FIPS Logs at the INFO level.
- **Centralized Management.** For clustered appliances, FIPS mode can only be turned on at the cluster level.
- **SSL Ciphers.** Only the following SSL ciphers are supported in FIPS mode: AES256-SHA:AES128-SHA:DES-CBC3-SHA.

Switching the Appliance to FIPS Mode

Use the `fipsconfig` CLI command to switch the appliance over to FIPS mode.



Note

Only administrators can use this command. A reboot is required after switching the appliance from non-FIPS mode to FIPS mode.

Before You Begin

Make sure that the appliance do not have any objects that are not FIPS compliant, for example, a DKIM verification profile with a key size of 512 bits. To enable FIPS mode, you must modify all the non-FIPS-compliant objects to meet FIPS requirements. See [Configuration Changes in FIPS Mode, page 27-1](#). For instructions to check if your appliance contains non-FIPS-compliant objects, see [Checking FIPS Mode Compliance, page 27-4](#).

Procedure

```
mail.example.com> fipsconfig
```

FIPS mode is currently disabled.

Choose the operation you want to perform:

- SETUP - Configure FIPS mode.
 - FIPSCHECK - Check for FIPS mode compliance.
- ```
[]> setup
```

To finalize FIPS mode, the appliance will reboot immediately. No commit will be required.

Are you sure you want to enable FIPS mode and reboot now ? [N]> y

Do you want to enable encryption of sensitive data in configuration file when FIPS mode is enabled? Changing the value will result in system reboot [N]> n

Enter the number of seconds to wait before forcibly closing connections.  
[30]>

System rebooting. Please wait while the queue is being closed...

Closing CLI connection.  
Rebooting the system...

## Encrypting Sensitive Data in FIPS Mode

Use the `fipsconfig` command to encrypt sensitive data such as passwords and keys, in your appliance. If you enable this option,

- The following critical security parameters in your appliance are encrypted and stored:
  - Certificate private keys
  - RADIUS passwords
  - LDAP bind passwords
  - Local users' password hashes
  - SNMP password
  - DK/DKIM signing keys
  - Outgoing SMTP authentication passwords
  - PostX encryption keys
  - PostX encryption proxy password
  - FTP Push log subscriptions' passwords
  - IPMI LAN password
  - Updater server URLs

**Note**

All users, including the administrators, cannot view the sensitive information in the configuration files.

- Swap space in your appliance is encrypted to prevent any unauthorized access or forensic attacks, if the physical security of the appliance is compromised.



**Procedure**

```
mail.example.com> fipsconfig
```

FIPS mode is currently enabled.

Choose the operation you want to perform:

```
- SETUP - Configure FIPS mode.
- FIPSCHECK - Check for FIPS mode compliance.
[]> setup
```

To finalize FIPS mode, the appliance will reboot immediately. No commit will be required.

Are you sure you want to disable FIPS mode and reboot now ? [N]> n

Do you want to enable encryption of sensitive data in configuration file when FIPS mode is enabled? Changing the value will result in system reboot [N]> y

Enter the number of seconds to wait before forcibly closing connections.  
[30]>

System rebooting. Please wait while the queue is being closed...

Closing CLI connection.  
Rebooting the system...

## Checking FIPS Mode Compliance

Use the `fipsconfig` command to check if your appliance contains any non-FIPS-compliant objects.

**Procedure**

```
mail.example.com> fipsconfig
```

FIPS mode is currently disabled.

Choose the operation you want to perform:

```
- SETUP - Configure FIPS mode.
- FIPSCHECK - Check for FIPS mode compliance.
[]> fipscheck
```

All objects in the current configuration are FIPS compliant.

FIPS mode is currently disabled.

## Managing Certificates and Keys

AsyncOS allows you to encrypt communications between the appliance and external machines by using a certificate and private key pair. You can upload an existing certificate and key pair, generate a self-signed certificate, or generate a Certificate Signing Request (CSR) to submit to a certificate authority to obtain a public certificate. The certificate authority will return a trusted public certificate signed by a private key that you can then upload onto the appliance.

When the appliance is in FIPS mode, you can continue to

The appliance's FIPS mode adds a number of restrictions to the certificates that the appliance uses in order for the appliance to be FIPS compliant. Certificates must use one of the following signature algorithms: SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512.

The appliance will not import certificates that do not use one of these algorithms. It also cannot be switched to FIPS mode if it has any non-compliant certificates in use on a listener. It will display an error message instead.

A `Non-FIPS` status for a certificate will be displayed in both the CLI and the GUI when the appliance is in FIPS mode. When selecting a certificate to use for a feature, such as a listener or destination control, the appliance does not display non-compliant certificates as an option.

See [Obtaining Certificates, page 23-2](#) for more information on using certificates on your appliance.

You can use FIPS-compliant certificates with any of the following services:

- **SMTP receiving and delivery.** Use the **Network > Listeners** page (or the `listenerconfig -> edit -> certificate` CLI command) to assign the certificate to any listeners that require encryption using TLS. You may want to only enable TLS on listeners facing the Internet (that is, public listeners), or you may want to enable encryption for all listeners, including internal systems (that is, private listeners).
- **Destination controls.** Use the **Mail Policies > Destination Controls** page (or the `destconfig` CLI command) to assign the certificate as a global setting to for all outgoing TLS connections for email delivery.
- **Interfaces.** Use the **Network > IP Interfaces** page (or the `interfaceconfig` CLI command) to enable the certificate for HTTPS services on an interface, including the management interface.
- **LDAP.** Use the **System Administration > LDAP** page to assign the certificate for all LDAP traffic that requires TLS connections. The appliance can also use LDAP for external authentication of users.

## Managing Keys for DKIM Signing and Verification

For an overview of how DomainKeys and DKIM work on the Email Security appliance, see [Chapter 20, “Email Authentication”](#).

### Related Topics

- [DKIM Signing, page 27-5](#)
- [DKIM Verification, page 27-6](#)

## DKIM Signing

When creating a DKIM signing key, you specify a key size. Email Security appliances in FIPS mode only support 2048 bits key size. The larger key size is more secure; however, larger keys can have an impact on performance.

The appliance cannot be switched to FIPS mode if it has any non-compliant RSA keys in use. It will display an error message instead.

FIPS-compliant signing keys are available for use in domain profiles and appear in the Signing Key list when creating or editing a domain profile using the **Mail Policies > Domain Profiles** page. Once you have associated a signing key with a domain profile, you can create a DNS text record which contains your public key. You do this via the **Generate** link in the DNS Text Record column in the domain profile listing (or via `domainkeysconfig -> profiles -> dnstxt` in the CLI).

## DKIM Verification

The appliance requires a message to use a FIPS-compliant key in order to verify a DKIM signature. If the signature does not use a FIPS-compliant key, the appliance returns a permanent failure.



## Using Email Security Monitor

---

- [Email Security Monitor Overview, page 28-1](#)
- [Email Security Monitor Pages, page 28-2](#)
- [Reporting Overview, page 28-35](#)
- [Managing Reports, page 28-36](#)
- [Troubleshooting Email Reports, page 28-39](#)

### Email Security Monitor Overview

The Email Security Monitor feature collects data from every step in the email delivery process. The database identifies and records each email sender by IP address, while interfacing with the SenderBase Reputation Service for real-time identity information. You can instantly report on any email sender's local mail flow history and show a profile that includes the sender's global record on the Internet. The Email Security Monitor feature allows your security team to “close the loop” on who is sending mail to your users, the amount of mail sent from and received by your users, and the effectiveness of your security policies.

This chapter explains how to:

- Access the Email Security Monitor feature to monitor inbound and outbound message flow.
- Make mail flow policy decisions (update whitelists, blacklists, and greylists) by querying for a sender's SenderBase Reputation Score (SBRS). You can query on network owners, domains, and even individual IP addresses.
- Report on mail flow, system status, and mail sent to and from your network.

For any given email sender for incoming mail, the Email Security Monitor database captures critical parameters such as:

- Message volume
- Connection history
- Accepted vs. rejected connections
- Acceptance rates and throttle limits
- Sender reputation filter matches
- Number of anti-spam messages for suspected spam and positively identified spam
- Number of virus-positive message detected by anti-virus scanning

See [Chapter 13, “Anti-Spam”](#) for more information on Anti-Spam scanning and [Chapter 12, “Anti-Virus”](#) for more information on anti-virus scanning.

The Email Security Monitor feature also captures information on which content filter a particular message triggers, including the internal user (email recipient) to or from which the message was sent.

The Email Security Monitor feature is available in the GUI only, and provides a view into your email traffic and the status of your appliance (including quarantines, work queues, and outbreaks). The appliance identifies when a sender falls outside of the normal traffic profile. Senders that do are highlighted in the interface, allowing you to take corrective action by assigning that sender to a sender group or refining the access profile of the sender; or, you can let AsyncOS’s security services continue to react and respond. Outbound mail has a similar monitoring capability, providing you a view into the top domains in the mail queue and the status of receiving hosts (see [Delivery Status Details Page, page 28-16](#)).

**Note**

Information for messages present in the work queue when the appliance is rebooted is not reported by the Email Security Monitor feature.

**Related Topics**

- [Email Security Monitor and Centralized Management, page 28-2](#)

## Email Security Monitor and Centralized Management

To view aggregated report data, deploy a Cisco Content Security Management appliance.

You cannot aggregate Email Security Monitor reports of clustered appliances. All reports are restricted to machine level. This means they cannot be run at the group or cluster levels — only on individual machines.

The same is true of the Archived Reports page — each machine in effect has its own archive. Thus, the “Generate Report” feature runs on the selected machine.

The Scheduled Reports page is not restricted to machine level; therefore, settings can be shared across multiple machines. Individual scheduled reports run at machine level just like interactive reports, so if you configure your scheduled reports at cluster level, every machine in the cluster will send its own report.

The “Preview This Report” button always runs against the login-host.

## Email Security Monitor Pages

The Email Security Monitor feature is comprised of all the pages available on the Monitor menu except the Quarantines pages.

You use these pages in the GUI to monitor domains that are connecting to the appliance’s listeners. 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 (including important information such as SBRS score and most recent sender group match for domains).

These pages help you classify mail relative to the appliance, and also relative to the services that exist beyond the scope of the gateway, such as the SenderBase Reputation Service, the Anti-Spam scanning service, the Anti-Virus scanning security services, content filters, and Outbreak Filters.

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. For information about generating PDFs in languages other than English, see the “[Notes on Reports](#)” section on page 28-36.

You can export graphs and other data to CSV (comma separated values) format via the **Export** link.

The exported CSV data will display all message tracking and reporting data in GMT regardless of what is set on the Email Security appliance. The purpose of the GMT time conversion is to allow data to be used independently from the appliance or when referencing data from appliances in multiple time zones.

**Note**

If you export localized CSV data, the headings may not render properly in some browsers. This occurs because some browsers may not use the correct character set for the localized text. To work around this problem, you can save the file to disk, and open the file using File > Open. When you open the file, select the character set to display the localized text.

For more information about automating the export of report data, see [Retrieving CSV Data, page 28-33](#)).

**Related Topics**

- [Searching and Email Security Monitor, page 28-4](#)
- [Viewing Details of Messages Included in Reports, page 28-4](#)
- [My Reports Page, page 28-5](#)
- [Overview Page, page 28-6](#)
- [Incoming Mail Page, page 28-9](#)
- [Outgoing Destinations, page 28-14](#)
- [Outgoing Senders, page 28-15](#)
- [Delivery Status Page, page 28-15](#)
- [Internal Users Page, page 28-16](#)
- [DLP Incidents Page, page 28-17](#)
- [Content Filters Page, page 28-18](#)
- [DMARC Verification Page, page 28-19](#)
- [Outbreak Filters Page, page 28-19](#)
- [Virus Types Page, page 28-21](#)
- [URL Filtering Page, page 28-21](#)
- [File Reputation and File Analysis Reports, page 28-22](#)
- [TLS Connections Page, page 28-22](#)
- [Inbound SMTP Authentication Page, page 28-23](#)
- [Rate Limits Page, page 28-23](#)
- [System Capacity Page, page 28-24](#)
- [System Status Page, page 28-30](#)
- [High Volume Mail Page, page 28-32](#)

- [Message Filters Page, page 28-32](#)
- [Retrieving CSV Data, page 28-33](#)

## 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 (IPv4 and IPv6)
- 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 IPv4 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).

For IPv6 address searches, AsyncOS supports the following formats:

- 2001:db8:2004:4202::0-2001:db8:2004:4202::ff
- 2001:db8:2004:4202::
- 2001:db8:2004:4202::23
- 2001:db8:2004:4202::/64

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

## Viewing Details of Messages Included in Reports

### Procedure

- 
- |               |                                                                                                                                                                         |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Click any blue number in a table on a report page.<br>(Not all tables have these links.)<br><br>The messages included in that number are displayed in Message Tracking. |
| <b>Step 2</b> | Scroll down to see the list.                                                                                                                                            |
- 

### Related Topics

- [Working with Message Tracking Search Results, page 29-4](#)

## My Reports Page

You can create a custom report page by assembling charts (graphs) and tables from existing report pages.

| To                                           | Do This                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Add modules to your custom report page       | <ol style="list-style-type: none"> <li>1. Go to <b>Monitor &gt; My Reports</b> and delete any sample modules that you do not need by clicking the [X] in the top right corner of the module.</li> <li>2. Do one of the following: <ul style="list-style-type: none"> <li>– Click the [+] button on a module in a report page under the Monitor menu to add it to your custom report.</li> <li>– Go to <b>Monitor &gt; My Reports</b>, click the [+] button in one of the sections, then select the report module that you want to add. You may need to check the + <b>Report Module</b> in each section to find the report that you are looking for.</li> </ul> </li> <li>3. Modules are added with default settings. If you add a module that you have customized (for example, by adding, deleting, or reordering columns), customize these modules again after adding them. Time range of the original module is not maintained.</li> <li>4. If you add a chart that includes a separate legend (for example, a graph from the Overview page), add the legend separately. If necessary, drag and drop it into position beside the data it describes.</li> </ol> <p>Notes:</p> <ul style="list-style-type: none"> <li>• Some modules on some report pages are available only using one of the above methods. If you cannot add a module using one method, try the other method.</li> <li>• You cannot add the following reporting modules to a custom report: <ul style="list-style-type: none"> <li>– The <b>Past Year Virus Outbreak Summary</b> chart and <b>Past Year Virus Outbreaks</b> table on the Outbreak Filters report page</li> <li>– Search results for all reports</li> </ul> </li> <li>• You can add each module only once; if you have already added a particular module to your report, the option to add it will not be available.</li> </ul> |
| View your custom report page                 | <ol style="list-style-type: none"> <li>1. Choose <b>Monitor &gt; My Reports</b>.</li> <li>2. For reports in the Time Range section: The time range selected for all report pages applies to all modules on the My Reports page. Select the time range to view.</li> </ol> <p>Newly-added modules appear at the top of the relevant section.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Rearrange modules on your custom report page | Drag and drop modules into the desired location.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Delete modules from your custom report page  | Click the [X] in the top right corner of the module.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |



## Overview Page

The Overview page provides a synopsis of the message activity of your appliance, including an overview of your quarantines and Outbreak Filters status (in the System Overview section of the page). The Overview page also includes graphs and detailed message counts for incoming and outgoing messages. You can use this page to monitor the flow of all mail into and out of your gateway.

The Overview page highlights how the appliance is integrated with the SenderBase Reputation Service for incoming mail (messages stopped by reputation filtering, for example). 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 sender reputation filtering (SBRS), messages with invalid recipients, messages marked as spam, messages marked as virus positive, and clean messages, over time.
- View the summary of the system status and local quarantines.
- See current virus and non-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.

### Related Topics

- [System Overview, page 28-6](#)
- [Incoming and Outgoing Summary and Graph, page 28-7](#)
- [Categorizing Email, page 28-8](#)
- [How Messages are Categorized, page 28-9](#)

## 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 outbreak activity.

### Related Topics

- [Status, page 28-6](#)
- [System Quarantines, page 28-7](#)
- [Virus Threat Level, page 28-7](#)

## 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 [Chapter 34, “Managing and Monitoring Using the CLI”](#) for more information.

**Incoming Messages:** The average rate of incoming mail per hour.

**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 Outbreak status as reported by the Threat Operations Center (TOC). 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 Outbreak Filters feature on your appliance.



### Note

In order for the Threat Level indicator to function, you need to have port 80 open on your firewall to “downloads.ironport.com.” Alternatively, if you have specified a local update server, the Threat Level indicator will attempt to use that address. The Threat Level indicator will also update correctly if you have configured a proxy for downloads via the Service Updates page. For more information, see [Service Updates, page 33-17](#).

Click the Outbreak Details link to view the external Threat Operations Center web site. Note that in order for this link to work, your appliance must be able to access the Internet. Note that the Separate Window icon ( ) indicates that a link will open in a separate window when clicked. You may need to configure your browser’s pop-up blocker settings to allow these windows.

## 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. The time range you select is used throughout all of the Email Security Monitor pages. The explanations of each type or category of message are below (see [Categorizing Email, page 28-8](#)).

While the mail trend graph displays a visual representation of the mail flow, the summary table provides a numeric breakdown of the same information. The summary table includes the percentage and actual number of each type of message, including the total number of attempted, threat, and clean messages.

The outgoing graph and summary show similar information for outbound mail.

### Related Topics

- [Notes on Counting Messages in Email Security Monitor, page 28-8](#)

## Notes on Counting Messages in Email Security Monitor

The method Email Security Monitor uses to count incoming mail depends on the number of recipients per message. For example, an incoming message from example.com sent to three recipients would count as three messages coming from that sender.

Because messages blocked by sender reputation filtering do not actually enter the work queue, the appliance does not have access to the list of recipients for an incoming message. In this case, a multiplier is used to estimate the number of recipients. This multiplier was determined by Cisco and based upon research of a large sampling of existing customer data.

## 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 (see [Notes on Counting Messages in Email Security Monitor, page 28-8](#)) 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.

**Detected by Advanced Malware Protection:** A message attachment was found to be malicious by file reputation filtering. This value does not include verdict updates or files found to be malicious by file analysis.

**Messages with Malicious URLs:** One or more URLs in the message were found to be malicious by URL filtering.

**Stopped by Content Filter:** The total count of messages that were stopped by a content filter.

**Stopped by DMARC:** The total count of messages that were stopped after DMARC verification.

**S/MIME Verification/Decryption Failed:** The total count of messages that failed S/MIME verification, decryption, or both.

**Marketing Messages:** The total count of marketing messages from legitimate sources, as determined by anti-spam scanning. This item appears only if marketing data are present in the system.

**S/MIME Verification/Decryption Successful:** The total count of messages that were successfully verified, decrypted, or decrypted and verified using S/MIME.

**Clean Messages:** 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.

## How Messages are Categorized

As messages proceed through the email pipeline, they can apply to multiple categories. For example, a message can be marked as spam or virus positive, it can also match a content filter. The various verdicts follow these rules of precedence: Outbreak Filters quarantining (in this case the message is not counted until it is released from the quarantine and again processed through the work queue), followed by spam positive, virus positive, and matching a content filter.

For example, if a message is marked as spam positive, and your anti-spam settings are set to drop spam positive messages, the message is dropped and the spam counter is incremented. Further, if your anti-spam settings are set to let the spam positive message continue on in the pipeline, and a subsequent content filter drops, bounces, or quarantines the message, the spam count is still incremented. The content filter count is only incremented if the message is not spam or virus positive.

## Incoming Mail Page

The **Incoming Mail** page provides a mechanism to report on the real-time information being collected by the Email Security Monitor feature for all remote hosts connecting to your appliance. This allows you to gather more information about an IP address, domain, and organization (network owner) sending mail to you. You can perform a Sender Profile search on IP addresses, domains, or organizations that have sent mail to you.

The Incoming Mail page has three views: Domain, IP Address, and Network Owner and provides a snapshot of the remote hosts connecting to the system in the context of the selected view.

It displays a table (Incoming Mail Details) of the top domains (or IP addresses, or network owners, depending on the view) that have sent mail to all public listeners configured on the appliance. You can monitor the flow of all mail into your gateway. You can click on any domain/IP/network owner to drill down to access details about this sender on a Sender Profile page (this is an Incoming Mail page, specific to the domain/IP/network owner you clicked on).

Not all available columns are displayed by default. You can show a different set of information by clicking the Columns link below the table. For example, you can show the "Detected by Advanced Malware Protection" column, which is hidden by default.

The Incoming Mail page extends to include a group of pages (Incoming Mail, Sender Profiles, and the Sender Group Report). From the **Incoming Mail** pages, you can:

- Perform a search on IP addresses, domains, or organizations (network owners) that have sent mail to you.
- View the Sender Groups report to see connections via a specific sender group and mail flow policy actions. See [Sender Groups Report, page 28-14](#) for more information.
- See detailed statistics on senders which have sent mail to you, including the number of attempted messages broken down by security service (sender reputation filtering, anti-spam, anti-virus, etc.).
- Sort by senders who have sent you a high volume of spam or virus email, as determined by anti-spam or anti-virus security services.
- Use the SenderBase Reputation service to drill down on and examine the relationship between specific IP addresses, domains, and organizations to obtain more information about a sender.

- Drill down on specific senders to obtain more information about a sender from the SenderBase Reputation Service, including a sender's SenderBase Reputation Score and which sender group the domain matched most recently. Add senders to sender groups.
- Drill down on a specific sender who sent a high volume of spam or virus email, as determined by the anti-spam or anti-virus security services.
- Once you have gathered information on a domain, you can add the IP address, domain, or organization to an existing sender group (if necessary) by clicking "Add to Sender Group" from a domain, IP address, or network owner profile page. See [Configuring the Gateway to Receive Email, page 5-1](#).

#### Related Topics

- [Incoming Mail, page 28-10](#)
- [Incoming Mail Details Listing, page 28-11](#)
- [Reporting Pages Populated with Data: Sender Profile Pages, page 28-12](#)
- [Sender Groups Report, page 28-14](#)

## 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.

See [Incoming Mail Details Listing, page 28-11](#) for an explanation of the data included in the Incoming Mail Details listing.

#### Related Topics

[Notes on Time Ranges in the Mail Trend Graph, page 28-10](#)

### 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 specify the time range to include in the results shown. Because the data is monitored in real time, information is periodically updated and summarized in the database.

Choose from the time range options in [Table 28-1](#).

**Table 28-1 Time Ranges Available in the Email Security Monitor Feature**

| This time range selected in the GUI | ...is defined as:                                       |
|-------------------------------------|---------------------------------------------------------|
| Hour                                | the last 60 minutes + up to 5 minutes                   |
| Day                                 | the last 24 hours + the last 60 minutes                 |
| Week                                | the last 7 days + the elapsed hours of the current day  |
| 30 days                             | the last 30 days + the elapsed hours of the current day |
| 90 days                             | the last 90 days + the elapsed hours of the current day |
| Yesterday                           | 00:00 to 23:59 (midnight to 11:59 PM)                   |

**Table 28-1 Time Ranges Available in the Email Security Monitor Feature (continued)**

| This time range selected in the GUI | ...is defined as:                                                                        |
|-------------------------------------|------------------------------------------------------------------------------------------|
| Previous Calendar Month             | 00:00 of the first day of the month to 23:59 of the last day of the month                |
| Custom Range                        | the range enclosed by the start date and hour and the end date and hour that you specify |

The time range options that you see will differ if you have enabled Centralized Reporting. For details, see information about Centralized Reporting Mode in [Chapter 42, “Centralizing Services on a Cisco Content Security Management Appliance.”](#)

## 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. See [Categorizing Email, page 28-8](#) for an explanation of the various categories.

The system acquires and verifies the validity of the remote host’s IP address (that is, the domain) by performing a *double DNS lookup*. For more information about double DNS lookups and sender verification, see [Configuring the Gateway to Receive Email, page 5-1](#).

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

The default Sender Detail 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).

The value for Stopped by Reputation Filtering is calculated based on several factors:

- Number of “throttled” messages from this sender.
- Number of rejected or TCP refused connections (may be a partial count).
- A conservative multiplier for the number of messages per connection.

When the appliance is under heavy load, an exact count of rejected connections is not maintained on a per-sender basis. Instead, rejected connections counts are maintained only for the most significant senders in each time interval. In this situation, the value shown can be interpreted as a “floor”; in other words, at least this many messages were stopped.



### Note

The Stopped by Reputation Filtering total on the Overview page is always based on a complete count of all rejected connections. Only the per-sender connection counts are ever limited due to load.

Additional columns that you can display are:

**Connections Rejected:** All connections blocked by HAT policies. When the appliance is under heavy load, an exact count of rejected connections is not maintained on a per-sender basis. Instead, rejected connections counts are maintained only for the most significant senders in each time interval.

**Connections Accepted:** All connections accepted

**Stopped by Recipient Throttling:** This is a component of Stopped by Reputation Filtering. It represents the number of recipient messages stopped because any of the following HAT limits have been exceeded: maximum recipients per hour, maximum recipients per message, or maximum messages per connection. This is summed with an estimate of the recipient messages associated with rejected or TCP refused connections to yield Stopped by Reputation Filtering.

**Detected by Advanced Malware Protection:** Messages with attachments that were found to be malicious by file reputation filtering. This value does not include verdict updates or files found to be malicious by file analysis.

**Total Threat:** Total number of threat messages (stopped by sender reputation, stopped as invalid recipient, spam, plus virus).

Show or hide columns by clicking the Column link at the bottom of the table.

Sort the listing by clicking the column header links. A small triangle beside the column header indicates the column by which the data is currently sorted.

#### Related Topics

- [“No Domain Information”, page 28-12](#)
- [Querying for More Information, page 28-12](#)

### “No Domain Information”

Domains which have connected to the appliance and could not be verified with a double-DNS lookup are automatically grouped into the special domain “No Domain Information.” You can control how these types of unverified hosts are managed via Sender Verification. See [Configuring the Gateway to Receive Email, page 5-1](#).

You can select the number of senders to show in the listing via the Items Displayed menu.

### Querying for More Information

For senders listed in the Email Security Monitor table, click the sender (or “No Domain Information” link) to drill down for more information on the particular sender. The results are displayed on a Sender Profile page which includes real-time information from the SenderBase Reputation Service. From the Sender Profile page, you can drill down for more information on specific IP addresses or network owners (see [Reporting Pages Populated with Data: Sender Profile Pages, page 28-12](#)).

You can also view another report, the Sender Groups report, by clicking the Sender Groups report link at the bottom of the Incoming Mail page. For more information about Sender Groups reports, see [Sender Groups Report, page 28-14](#).

## Reporting Pages Populated with Data: Sender Profile Pages

If you clicked a sender in the Incoming Mail Details table on an 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. For more information on this relationship and how it relates to the SenderBase Reputation Service, see [Configuring the Gateway to Receive Email, page 5-1](#).

The Sender Profile pages displayed for IP addresses, network owners, and domains vary slightly. For each, the page contains a graph and summary table for incoming mail from this sender. Below the graph is a table listing domains or IP addresses associated with the sender (the Sender Profile page for individual IP addresses does not contain the detailed listing) and an information section with the current SenderBase, sender group, and network information for the sender.

- Network Owner profile pages contain information for the network owner, as well as the domains and IP addresses associated with that network owner.
- Domain profile pages contain information for the domains and IP addresses associated with that domain.
- IP address profile pages contain information about the IP address only.

Each sender profile page contains the following data in the Current Information table at the bottom of the page:

- The **Global** information from the SenderBase Reputation Service, including:
  - IP Address, Domain Name, and/or Network Owner
  - Network Owner Category (Network Owner Only)
  - CIDR Range (IP addresses only)
  - Daily Magnitude and Monthly Magnitude for the IP address, Domain, and/or Network Owner
  - Days since the first message was received from this sender
  - Last sender group and whether DNS verified (IP Address sender profile page only)

Daily magnitude is a measure of how many messages a domain has sent over the last 24 hours. Similar to the Richter scale used to measure earthquakes, SenderBase magnitude is a measure of message volume calculated using a log scale with a base of 10. The maximum theoretical value of the scale is set to 10, which equates to 100% of the world's email message volume (approximately 10 billion messages/day). Using the log scale, a one-point increase in magnitude equates to a 10x increase in actual volume.

Monthly magnitude is calculated using the same approach as daily magnitude, except the percentages are calculated based on the volume of email sent over the last 30 days.

- Average Magnitude (IP addresses only)
- Lifetime Volume / 30 Day Volume (IP address profile pages only)
- Bonded Sender Status (IP address profile pages only)
- SenderBase Reputation Score (IP address profile pages only)
- Days Since First Message (network owner and domain profile pages only)
- Number of Domains Associated with this Network Owner (network owner and domain profile pages only)
- Number of IP Addresses in this Network Owner (network owner and domain profile pages only)
- Number of IP Addresses used to Send Email (network owner pages only)

Click the “More from SenderBase” link to see a page with all information supplied by the SenderBase Reputation Service.

- The **Mail Flow Statistics** information, with Email Security Monitor information collected about the sender over the time range that you specify.
- **Details** about the domains and IP addresses controlled by this network owner are displayed on network owner profile pages. Details about the IP addresses in the domain are displayed on domain pages.



From a domain profile page, you can drill down to a specific IP address, or drill up to view an organization profile page. You can also display the DNS Verified status, SBRS (SenderBase Reputation Score), and Last Sender Group for each sender address in the IP Addresses table by clicking the Columns link at the bottom of that table. You can also hide any columns in that table.

From a network owner profile page, you can display information such as Connections Rejected, Connections Accepted, Stopped by Recipient Throttling, and Detected by Advanced Malware Protection for each domain in the Domains table by clicking the Columns link at the bottom of that table. You can also hide any columns in that table.

If you are an administrator of the system, on each of these pages, you can choose to add the network owner, domain, or IP address to a sender group by clicking the check box for the entity (if necessary) and then clicking Add to Sender Group.

You can also add a sender to a sender group by clicking the **Add to Sender Group** link below the Sender Group Information in the Current Information table for the sender and clicking Add to Sender Group. For more information about adding senders to sender groups, see [Configuring the Gateway to Receive Email, page 5-1](#). Of course, you do not have to make any changes — you can let the security services handle incoming mail.

#### Related Topics

- [Sender Profile Search, page 28-14](#)

## Sender Profile Search

Type an IP address, a domain, or an organization name in the Quick Search box to search for a specific sender.

A Sender Profile page is displayed with the information for sender. See [Reporting Pages Populated with Data: Sender Profile Pages, page 28-12](#).

## Sender Groups Report

The Sender Groups report provides a summary of connections by sender group and mail flow policy action, allowing you to review SMTP connection and mail flow policy trends. The Mail Flow by Sender Group listing shows the percentage and number of connections for each sender group. The Connections by Mail Flow Policy Action chart shows the percentage of connections for each mail flow policy action. This page provides an overview of the effectiveness of your Host Access Table (HAT) policies. For more information about the HAT, see [Configuring the Gateway to Receive Email, page 5-1](#).

## 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, such as an hour, a week, or a custom range. 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 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).



### Note

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, such as an hour, a week, or a custom range. 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?

## 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.

The **Delivery Status Page** displays the same information as the `tophosts` command within the CLI. (For more information, see “Determining the Make-up of the Email Queue” in [Chapter 34, “Managing and Monitoring Using the CLI.”](#))

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.

**Note**

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.

**Related Topics**

- [Retrying Delivery, page 28-16](#)
- [Delivery Status Details Page, page 28-16](#)

## Retrying Delivery

Messages that are scheduled for later delivery can be immediately retried by clicking **Retry All Delivery**. Retry All Delivery allows you to reschedule messages in the queue for immediate delivery. All domains that are marked as “down” and any scheduled or soft bounced messages are queued for immediate delivery.

To retry delivery to a specific destination domain, click the domain name link. On the Delivery Status Details page, click **Retry Delivery**.

You can also use the `delivernow` command in the CLI to reschedule messages for immediate delivery. For more information, see [Scheduling Email for Immediate Delivery, page 34-31](#).

## Delivery Status Details Page

Use the **Delivery Status Details Page** to look up statistics on a specific recipient domain. This page displays the same information as the `hoststatus` command within the CLI: Mail Status, Counters and Gauges. (For more information, see “Monitoring the Status of a Mail Host” in [Chapter 34, “Managing and Monitoring Using the CLI.”](#)) To search for a specific domain, type the name of the domain in the Domain Name: field and click **Search**. Virtual Gateway address information appears if you are using the `altsrchost` feature.

## Internal Users Page

The Internal Users page provides information about the mail sent and received by your internal users, *per email address* (a single user may have multiple email addresses listed — the email addresses are not combined in the report).

The page consists of two sections: graphs depicting the top users by clean incoming and outgoing messages, and user mail flow details. 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. You can also display hidden table columns or hide default columns by clicking the Columns link below the table.

The User Mail Flow Details listing breaks down the mail received and sent by each email address into Clean, Spam Detected (incoming only), Virus Detected, and Content Filter Matches. You can sort the listing by clicking on the column headers.

Using the Internal Users report, you can answer these kinds of questions:

- Who is sending the most external email?
- Who receives the most clean email?
- Who receives the most spam?

- Who is triggering which content filters?
- Whose email is getting caught by content filters?

Inbound Internal Users are the users for which you received email, based on the Rcpt To: address. Outbound Internal Users are based on the Mail From: address and are useful when tracking the types of email that senders on your internal network are sending.

Note that some outbound mail (like bounces) have a null sender. They are counted under outbound and “unknown.”

Click on an internal user to view the Internal User detail page for that user.

Click the Columns link below the table to show columns that are hidden by default, such as the Incoming Detected by Advanced Malware Protection column.

#### Related Topics

- [Internal User Details, page 28-17](#)
- [Searching for a Specific Internal User, page 28-17](#)

## 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). Optionally, for incoming messages, you can click the Columns link below the table to show the Incoming Detected by Advanced Malware Protection column. This value reflects the number messages that contained attachments that were determined by file reputation filtering to be malicious. It does not include verdict updates or files found to be malicious by file analysis. Incoming and outgoing content filter and DLP policy matches are also shown.

Click on a content filter name to view detailed information for that filter in the corresponding content filter information page (see [Content Filters Page, page 28-18](#)). You can use this method to get a list of users who also sent or received mail that matched that particular content filter.

## Searching for a Specific Internal User

You can search for a specific internal user (email address) via the search form at the bottom of the Internal Users page and the Internal User detail page. 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”).

## DLP Incidents Page

The DLP Incidents page shows information on the incidents of data loss prevention (DLP) policy violations occurring in outgoing mail. The appliance uses the DLP email policies enabled in the Outgoing Mail Policies table to detect sensitive data sent by your users. Every occurrence of an outgoing message violating a DLP policy is reported as an incident.

Using the DLP Incidents report, you can answer these kinds of questions:

- What type of sensitive data is being sent by your users?
- How severe are these DLP incidents?
- How many of these messages are being delivered?
- How many of these messages are being dropped?

- Who is sending these messages?

The DLP Incidents page is comprised of two main sections:

- the DLP incident trend graphs summarizing the top DLP incidents by severity (Low, Medium, High, Critical) and policy matches, and
- the DLP Incidents Details listing.

You can select a time range on which to report, such as an hour, a week, or a custom range. As with all reports, you can export the data for the graphs or the details listing to CSV format via the **Export** link or PDF format by clicking the **Printable (PDF)** link. For information about generating PDFs in languages other than English, see the “[Notes on Reports](#)” section on page 28-36.

Click on the name of a DLP policy to view detailed information on the DLP incidents detected by the policy. You can use this method to get a list of users who sent mail that contained sensitive data detected by the policy.

#### Related Topics

- [DLP Incidents Details, page 28-18](#)
- [DLP Policy Detail Page, page 28-18](#)

## DLP Incidents Details

The DLP policies currently enabled in the appliance’s outgoing mail policies are listed in the DLP Incidents Details table at the bottom of the DLP Incidents page. Click on the name of a DLP policy to view more detailed information.

The DLP Incidents Details table shows the total number of DLP incidents per policy, with a breakdown by severity level, and the number of messages delivered in the clear, delivered encrypted, or dropped. Click on the column headings to sort the data.

## DLP Policy Detail Page

If you clicked the name of a DLP policy in the DLP Incidents Details table, the resulting DLP Policy Detail page displays the DLP incidents data for the policy. The page displays graphs on the DLP incidents based on severity.

The page also includes an Incidents by Sender listing at the bottom of the page that lists each internal user who has sent a message that violated the DLP policy. The listing also shows the total number of DLP incidents for this policy per user, with a breakdown by severity level, and whether any of the messages were delivered in the clear, delivered encrypted, or dropped. You can use the Incidents by Sender listing to find out which users may be sending your organization’s sensitive data to people outside your network.

Clicking on the sender name opens up the Internal Users page. See [Internal Users Page, page 28-16](#) for more information.

## Content Filters Page

The Content Filters page shows information about the top incoming and outgoing 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.

#### Related Topics

- [Content Filter Details, page 28-19](#)

## Content Filter Details

The Content Filter detail page displays matches for that filter over time, as well as matches by internal user.

In the Matches by Internal User section, you can click the name of a user to view that internal user's (email address) Internal User details page (see [Internal User Details, page 28-17](#)).

## DMARC Verification Page

The DMARC Verification page shows the top domains that failed DMARC verification and the details of actions AsyncOS performed on the messages that failed DMARC verification. You can use this report to fine-tune your DMARC settings and answer these kinds of questions:

- Which are the domains that sent maximum number of messages that are not DMARC compliant?
- For each domain, what are the actions AsyncOS performed on the messages that failed DMARC verification?

The DMARC Verification page contains:

- A bar chart showing top domains by DMARC verification failures.
- Tabular representation of the following, for each domain:
  - Number of messages that were rejected, quarantined, or accepted without taking any action. Click on the number to view a list of messages under the selected category.
  - Number messages that passed DMARC verification.
  - Total number of DMARC verification attempts.

You can select a time range on which to report, such as an hour, a week, or a custom range. As with all reports, you can export the data for the graphs or the details listing to CSV format via the **Export** link or PDF format by clicking the **Printable (PDF)** link.

## Outbreak Filters Page

The Outbreak Filters page shows the current status and configuration of Outbreak Filters on your appliance as well as information about recent outbreaks and messages quarantined due to Outbreak Filters. You can use this page to monitor your defense against targeted virus, scam, and phishing attacks.

The Threats By Type section shows the different types of threat messages received by the appliance.

The Threat Summary section shows a breakdown of the threat messages by Malware, Phish, Scam, and Virus. Click on the number to view a list of all the messages that are included in that number using Message Tracking.

The Past Year Outbreak Summary lists global as well as local outbreaks over the past year, allowing you to compare local network trends to global trends. The listing of global outbreaks is a superset of all outbreaks, both viral and non-viral, whereas local outbreaks are limited to virus outbreaks that have affected your appliance. Local outbreak data does not include non-viral threats. Global outbreak data represents all outbreaks detected by the Threat Operations Center which exceeded the currently configured threshold for the outbreak quarantine. Local outbreak data represents all virus outbreaks detected on this appliance which exceeded the currently configured threshold for the outbreak quarantine. The Total Local Protection Time is always based on the difference between when each virus outbreak was detected by the Threat Operations Center and the release of an anti-virus signature by a major vendor. Note that not every global outbreak affects your appliance. A value of "--" indicates either a protection time does not exist, or the signature times were not available from the anti-virus vendors (some vendors may not report signature times). This does not indicate a protection time of zero, rather it means that the information required to calculate the protection time is not available.

The Quarantined Messages section summarizes Outbreak Filters quarantining, and is a useful gauge of how many potential threat messages Outbreak Filters are catching. Quarantined messages are counted at time of release. Typically, messages will be quarantined before anti-virus and anti-spam rules are available. When released, they will be scanned by the anti-virus and anti-spam software and determined to be positive or clean. Because of the dynamic nature of Outbreak tracking, the rule under which a message is quarantined (and even the associated outbreak) may change while the message is in the quarantine. Counting the messages at the time of release (rather than the time of entry into the quarantine) avoids the confusion of having counts that increase and decrease.

The Threat Details listing displays information about specific outbreaks, including the threat category (virus, scam, or phishing), threat name, a description of the threat, and the number of messages identified. For virus outbreaks, the Past Year Virus Outbreaks include the Outbreak name and ID, time and date a virus outbreak was first seen globally, the protection time provided by Outbreak filters, and the number of quarantined messages. You can select either global or local outbreaks as well as the number of messages to display via the menu on the left. You can sort the listing by clicking on the column headers. Click on the number to view a list of all the messages that are included in that number using Message Tracking.

The First Seen Globally time is determined by the Threat Operations Center, based on data from SenderBase, the world's largest email and web traffic monitoring network. The Protection Time is based on the difference between when each threat was detected by the Threat Operations Center and the release of an anti-virus signature by a major vendor.

A value of "--" indicates either a protection time does not exist, or the signature times were not available from the anti-virus vendors (some vendors may not report signature times). This does not indicate a protection time of zero. Rather, it means that the information required to calculate the protection time is not available.

Hit Messages from Incoming Messages section shows the percentage and number of viral attachment, other threats (non-viral), and clean incoming messages.

Hit Messages by Threat Level section shows the percentage and number of incoming threat messages (viral and non-viral) based on threat levels (Level 1 through 5).

Messages resided in Outbreak Quarantine section shows the number of threat messages resided in the Outbreak Quarantine based on the duration.

Top URL's Rewritten section shows the list of top 10 URLs that were rewritten based on the number of occurrences. Use the Items Displayed drop-down to view more rewritten URLs. Click on the number to view a list of all the messages that contain the selected rewritten URL on the Message Tracking page.

Using the Outbreak Filters page, you can answer questions like:

- How many messages are being quarantined and what type of threats were they?
- How much lead time has the Outbreak Filter feature been providing for virus outbreaks?
- How do my local virus outbreaks compare to the global outbreaks?

## 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 appliance. 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.

If you run multiple virus scanning engines, the Virus Types page includes results from all enabled virus scanning engines. The name of the virus displayed on the page is a name determined by the virus scanning engines. If more than one scanning engine detects a virus, it is possible to have more than one entry for the same virus.

The Virus Types page gives you an overview of the viruses entering or being sent from or to your network. The Top Incoming Virus Detected section shows a chart view of the viruses that have been sent to your network in descending order. The Top Outgoing Virus Detected section shows a chart view of the viruses that have been sent from your network in descending order.



### Note

To see which hosts sent virus-infected messages to your network, you can go to the Incoming Mail page, specify the same reporting period and sort by virus-positive. Similarly, to see which IP addresses have sent virus-positive email within your network, you can view the Outgoing Senders page and sort by virus-positive messages.

The VirusTypes Details listing displays information about specific viruses, including the infected incoming and outgoing messages, and the total infected messages. The details listing for infected incoming messages displays the name of the virus and the number of incoming messages infected with this virus. Similarly, the outgoing messages displays the name of the virus and the number of outgoing messages infected with the virus. You can sort the Virus Type details by Incoming Messages, Outgoing Messages, or Total Infected Messages.

## URL Filtering Page

- URL Filtering report modules are populated only if URL filtering is enabled.
- URL Filtering reports are available for incoming and outgoing messages.
- Only messages that are scanned by the URL filtering engine (either as part of anti-spam/outbreak filter scanning or through message/content filters) are included in these modules. However, not all of the results are necessarily specifically attributable to the URL Filtering feature.
- The Top URL Categories module includes all categories found in messages that have been scanned, whether or not they match a content or message filter.
- Each message can be associated with only one URL reputation level. For messages with multiple URLs, the statistics reflect the lowest reputation of any URL in the message.
- URLs in the global whitelist configured at Security Services > URL Filtering are not included in reports.



URLs in whitelists used in individual filters are included in reports.

- Malicious URLs are URLs that Outbreak Filters have determined to have poor reputation. Suspicious URLs are those that Outbreak Filters have determined to require click-time protection. Suspicious URLs have therefore been rewritten to redirect them to the Cisco Web Security proxy.
- Results of URL category-based filters are reflected in content and message filter reports.
- Results of click-time URL evaluations by the Cisco Web Security proxy are not reflected in reports.

## File Reputation and File Analysis Reports

For the following reports, see [File Reputation and File Analysis Reporting and Tracking, page 16-11](#):

- Advanced Malware Protection
- File Analysis
- AMP Verdict Updates

## 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 TLS Connections page is divided into a section for incoming connections and a section for outgoing connections. Each section includes a graph, summaries, and a table with details.

The graph displays a view of incoming or outgoing TLS-encrypted and non-encrypted connections over the time range you specify. The graph displays the total volume of messages, the volume of encrypted and unencrypted messages, and the volume of successful and failed TLS encrypted messages. The graphs distinguish between connections in which TLS was required and connections in which TLS was merely preferred.

The table displays details for domains sending or receiving encrypted messages. For each domain, you can view the number of required and preferred TLS connections that were successful and that failed, the total number of TLS connections attempted (whether successful or failed), and the total number of unencrypted connections. You can also view the percentage of all connections in which TLS was attempted, and the total number of encrypted messages sent successfully, regardless of whether TLS was preferred or required. You can show or hide columns by clicking the Columns link at the bottom of this table.

## Inbound SMTP Authentication Page

The Inbound SMTP Authentication page shows the use of client certificates and the SMTP AUTH command to authenticate SMTP sessions between the Email Security appliance and users' mail clients. If the appliance accepts the certificate or SMTP AUTH command, it will establish a TLS connection to the mail client, which the client will use to send a message. Since it is not possible for the appliance to track these attempts on a per-user basis, the report shows details on SMTP authentication based on the domain name and domain IP address.

Use this report to determine the following information:

- Overall, how many incoming connections use SMTP authentication?
- How many connections use a client certificate?
- How many connections use SMTP AUTH?
- What domains are failing to connect when attempting to use SMTP authentication?
- How many connections are successfully using the fall-back when SMTP authentication fails?

The Inbound SMTP Authentication page includes a graph for received connections, a graph for mail recipients who attempted an SMTP authentication connection, and a table with details on the attempts to authenticate connections.

The Received Connections graph shows the incoming connections from mail clients that attempt to authenticate their connections using SMTP authentication over the time range you specify. The graph displays the total number of connections the appliance received, the number that did not attempt to authenticate using SMTP authentication, the number that failed and succeeded to authenticate the connection using a client certificate, and the number that failed and succeeded to authenticate using the SMTP AUTH command.

The Received Recipients graph displays the number of recipients whose mail clients attempted to authenticate their connections to the Email Security appliances to send messages using SMTP authentication. The graph also shows the number of recipients whose connections were authenticated and the number of recipients whose connections were not authenticated.

The SMTP Authentication details table displays details for the domains whose users attempt to authenticate their connections to the Email Security appliance to send messages. For each domain, you can view the number of connection attempts using a client certificate that were successful or failed, the number of connection attempts using the SMTP AUTH command that were successful or failed, and the number that fell back to the SMTP AUTH after their client certificate connection attempt failed. You can use the links at the top of the page to display this information by domain name or domain IP address.

## Rate Limits Page

Rate Limiting by envelope sender allows you to limit the number of email message recipients per time interval from an individual sender, based on the mail-from address. The Rate Limits report shows you the senders who most egregiously exceed this limit.

Use this report to help you identify the following:

- Compromised user accounts that might be used to send spam in bulk.
- Out-of-control applications in your organization that use email for notifications, alerts, automated statements, etc.
- Sources of heavy email activity in your organization, for internal billing or resource-management purposes.

- Sources of large-volume inbound email traffic that might not otherwise be considered spam.

Note that other reports that include statistics for internal senders (such as Internal Users or Outgoing Senders) measure only the number of messages sent; they do not identify senders of a few messages to a large number of recipients.

The Top Offenders by Incident chart shows the envelope senders who most frequently attempted to send messages to more recipients than the configured limit. Each attempt is one incident. This chart aggregates incident counts from all listeners.

The Top Offenders by Rejected Recipients chart shows the envelope senders who sent messages to the largest number of recipients above the configured limit. This chart aggregates recipient counts from all listeners.

To configure rate limiting by envelope sender or modify the existing rate limit, see [Defining Rules for Incoming Messages Using a Mail Flow Policy, page 7-15](#).

## 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 a 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.

It is important to monitor your appliance to ensure that your capacity is appropriate to your message volumes. Over time, volume will inevitably rise and appropriate monitoring will ensure that additional capacity or configuration changes can be applied proactively. The most effective way to monitor system capacity is to track overall volume, messages in the work queue and incidents of Resource Conservation Mode.

- **Volume:** It is important to have an understanding of the “normal” message volume and the “usual” spikes in your environment. Track this data over time to measure volume growth. You can use the Incoming Mail and Outgoing Mail pages to track volume over time. For more information, see [System Capacity- Incoming Mail, page 28-26](#) and [System Capacity-Outgoing Mail, page 28-27](#).
- **Work Queue:** The work queue is designed to work as a “shock absorber”-- absorbing and filtering spam attacks and processing unusual increases in ham messages. However, the work queue is also the best indicator of a system under stress, prolonged and frequent work queue backups may indicate a capacity problem. You can use the WorkQueue page to track the average time messages spend in the work queue and the activity in your work queue. For more information, see [System Capacity-Workqueue, page 28-25](#).
- **Resource Conservation Mode:** When a appliance becomes overloaded, it will enter “Resource Conservation Mode” (RCM) and send a CRITICAL system alert. This is designed to protect the device and allow it to process any backlog of messages. Your appliance should enter RCM infrequently and only during a very large or unusual increase in mail volume. Frequent RCM alerts may be an indication that the system is becoming overloaded. Resource Conservation Mode is not tracked by the system capacity page.

**Related Topics**

- [System Capacity- Workqueue, page 28-25](#)
- [System Capacity- Incoming Mail, page 28-26](#)
- [System Capacity-Outgoing Mail, page 28-27](#)
- [System Capacity-System Load, page 28-28](#)
- [Note about Memory Page Swapping, page 28-29](#)
- [System Capacity- All, page 28-30](#)

## System Capacity- Workqueue

The Workqueue page shows the average time a message spends in the work queue, excluding any time spent in the Spam quarantine or in a policy, virus, or outbreak quarantine. You can view time periods from an hour up to one month. This average can help in identifying both short term events delaying mail delivery and identify long term trends in the workload on the system.

**Note**

---

If a message is released from the quarantine into the work queue, the “average time in work queue” metric ignores this time. This prevents double-counting and distorted statistics due to extended time spent in a quarantine.


---

The report also shows the volume of messages in the work queue over a specified time period, and it shows the maximum messages in the work queue over the same time period.

Occasional spikes in the Workqueue graphs are normal and expected. If the spikes occur with increasing frequency and are maintained over a long period of time, this may indicate a capacity issue. When reviewing the work queue page, you may want to measure the frequency of work queue backups, and take note of work queue backups that exceed 10,000 messages.

**Figure 28-1 System Capacity - Workqueue****System Capacity**

[ Workqueue | Incoming Mail | Outgoing Mail | System Load | All ]


[Printable \(PDF\)](#)

## System Capacity- Incoming Mail

The incoming mail page shows incoming connections, the total number of incoming messages, the average message size, and the total incoming message size. You can limit the results to the time range that you specify. It is important to have an understanding of the trends of normal message volume and spikes in your environment. You can use the incoming mail page to help track volume growth over time and plan for system capacity. You might also want to compare the Incoming Mail data with the Sender Profile data to view the trends in volumes of emails that are being sent from specific domains to your network.



**Note**

an increased number of incoming connections may not necessarily affect system load.

**Figure 28-2 System Capacity - Incoming Mail (Page 1 of 2)****Figure 28-3 System Capacity - Incoming Mail (Page 2 of 2)**

## System Capacity-Outgoing Mail

The outgoing mail page shows outgoing connections, the total number of outgoing messages, the average message size, and the total outgoing message size. You can limit the results to the time range that you specify. It is important to have an understanding of the trends of normal message volume and spikes in your environment. You can use the outgoing mail page to help track volume growth over time and plan for system capacity. You might also want to compare the Outgoing Mail data with the Outgoing Destinations data to view the trends in volumes of emails that are being sent from specific domains or IP addresses.

**Figure 28-4** System Capacity - Outgoing Mail (page 1 of 2)**Figure 28-5** System Capacity - Outgoing Mail (page 2 of 2)


## System Capacity-System Load

The system load report shows the overall CPU usage on your appliance. AsyncOS is optimized to use idle CPU resources to improve message throughput. High CPU usage may not indicate a system capacity problem. If the high CPU usage is coupled with consistent, high-volume memory page swapping, you may have a capacity problem. This page also shows a graph that displays the amount of CPU used by different functions, including mail processing, spam and virus engines, reporting, and quarantines. The CPU-by-function graph is a good indicator of which areas of the product use the most resources on your system. If you need to optimize your appliance, this graph can help you determine which functions may need to be tuned or disabled.

The memory page swapping graph shows how frequently the system must page to disk.

**Figure 28-6 System Capacity - System Load****System Capacity**[ Workqueue | Incoming Mail | Outgoing Mail | **System Load** | All ]


Printable (PDF)



## Note about Memory Page Swapping

The system is designed to swap memory regularly, so some memory swapping is expected and is not an indication of problems with your appliance. Unless the system *consistently* swaps memory in high volumes, memory swapping is normal and expected behavior (especially on C160 appliances). For example, [Figure 28-7](#) shows a system that consistently swaps memory in high volumes. To improve performance, you may need to add appliances to your network or tune your configuration to ensure maximum throughput.



**Figure 28-7 System Capacity - System Load (System Under Heavy Load)**

## System Capacity- All

The All page consolidates all the previous system capacity reports onto a single page so you can view the relationship between the different reports. For example, you might view the message queue is high at the same time that excessive memory swapping takes place. This might be an indication that you have a capacity problem. You may want to save this page as PDF to preserve a snapshot of system performance for later reference (or to share with support staff). For information about generating PDFs in languages other than English, see the [“Notes on Reports”](#) section on page 28-36.

## System Status Page

The **System Status** page provides a detailed representation of all real-time mail and DNS activity for the system. The information displayed is the same information that is available by using the `status detail` and `dnsstatus` commands in the CLI. For more information, see “Monitoring Detailed Email Status” for the `status detail` command and “Checking the DNS Status” for the `dnsstatus` command in [Chapter 34, “Managing and Monitoring Using the CLI.”](#)

The System Status page is comprised of four sections: System Status, Gauges, Rates, and Counters.

### Related Topics

- [System Status, page 28-30](#)
- [Gauges, page 28-31](#)
- [Rates, page 28-31](#)
- [Counters, page 28-32](#)

## System Status

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

**Related Topics**

- [Mail System Status, page 28-31](#)
- [Version Information, page 28-31](#)

**Mail System Status**

The Mail System Status section includes:

- System Status (for more information about system status, see [Status, page 28-6](#))
- 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 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.

This information is useful if you are contacting Cisco Customer Support. (See [Working with Technical Support, page 40-28](#).)

**Gauges**

The Gauges section shows queue and resource utilization.

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

Mail Gateway Appliance refers to the percentage of the CPU that AsyncOS processes are consuming. CASE refers to several items, including the Anti-Spam scanning engine and Outbreak Filters processes.

- General Resource Utilization
- Logging Disk 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.

**Note**

Only user accounts that are in the administrator or operator group have access to reset the counters. User accounts you create in the guest group will not be able to reset the counters. For more information, see [Working with User Accounts, page 32-1](#).

Click Reset Counters to reset the counters. This button offers the same functionality as the `resetcounters` command in the CLI. For more information, see [Resetting Email Monitoring Counters, page 34-21](#).

- Mail Handling Events
- Completion Events
- Domain Key Events
- DNS Status

## High Volume Mail Page

**Note**

The High Volume Mail page shows data only from message filters that use Header Repeats rule.

The High Volume Mail page contains the following reports in the form of bar charts:

- **Top Subjects.** You can use this chart to understand the top subjects of messages that AsyncOS received.
- **Top Envelope Senders.** You can use this chart to understand the top envelope senders of messages that AsyncOS received.
- **Top Message Filters by Number of Matches.** You can use this chart to understand the top message filter (that uses Header Repeats rule) matches.

The High Volume Mail page also provides a tabular representation of the top message filters and the number of matches for the respective message filters. Click on the number to view a list of all the messages that are included in that number using Message Tracking.

You can select a time range on which to report, such as an hour, a week, or a custom range. As with all reports, you can export the data for the graphs or the details listing to CSV format via the **Export** link or PDF format by clicking the **Printable (PDF)** link.

## Message Filters Page

The Message Filters page shows information about the top message filter matches (which message filter had the most matching messages) in two forms: a bar chart and a tabular representation.

Using the bar chart, you can find the message filters that are being triggered the most by incoming and outgoing messages. The tabular representation shows the top message filters and the number of matches for the respective message filters. Click on the number to view a list of all the messages that are included in that number using Message Tracking.

You can select a time range on which to report, such as an hour, a week, or a custom range. As with all reports, you can export the data for the graphs or the details listing to CSV format via the **Export** link or PDF format by clicking the **Printable (PDF)** link.

## Retrieving CSV Data

You can retrieve the data used to build the charts and graphs in the Email Security Monitor in CSV format. The CSV data can be accessed in two ways:

- **CSV reports delivered via email.** You can generate a CSV report that is delivered via email or archived. This delivery method is useful when you want separate reports for each table represented on an Email Security Monitor page, or when you want to send CSV data to users who do not have access to internal networks.

The comma-separated values (CSV) Report Type is an ASCII text file which contains the tabular data of the scheduled report. Each CSV file may contain up to 100 rows. If a report contains more than one type of table, a separate CSV file will be created for each table. Multiple CSV files for a single report will be compressed into a single .zip file for the archived file storage option or will all be attached to separate e-mail messages for e-mail delivery.

For information about configuring scheduled or on-demand reports, see [Reporting Overview, page 28-35](#).

- **CSV files retrieved via HTTP.** You can retrieve the data used to build the charts and graphs in the Email Security Monitor feature via HTTP. This delivery method is useful if you plan to perform further analysis on the data via other tools. You can automate the retrieval of this data, for example, by an automatic script that will download raw data, process, and then display the results in some other system.

### Related Topics

- [Retrieving CSV Data Via Automated Processes, page 28-33](#)

## Retrieving CSV Data Via Automated Processes

The easiest way to get the HTTP query you will need is to configure one of the Email Security Monitor pages to display the type of data you want. You can then copy the **Export** link. This is the download URL. When automating data retrieval like this it is important to note which parameters in the download URL should be fixed and which should change (see below).

The download URL is encoded in such a way that it can be copied to an external script that can execute the same query (using proper HTTP authentication) and get a similar data set. The script can use Basic HTTP Authentication or cookie authentication. Keep the following in mind when retrieving CSV data via automated processes:

- Time range selection (past hour, day, week, etc) in relation to when the URL is used again. If you copy the URL to retrieve a CSV data set for “Past Day,” the next time you use that URL you will get a new data set that covers the “Past Day” from the time you send the URL again. The date range selection is retained, and appears in the CSV query string (e.g. `date_range=current_day`).
- Filtering and grouping preferences for the data set. Filters are retained and appear in the query string. Note that filters in reports are rare — one example is the “Global / Local” outbreaks selector in the Outbreaks report.
- The CSV download returns all rows of data in the table for the selected time range.

- The CSV download returns the rows of data in the table ordered by timestamp and key. You can perform further sorting in a separate step such as via a spreadsheet application.
- The first row contains column headers that match the display names shown in the report. Note that timestamps (see [Timestamps, page 28-34](#)) and keys (see [Keys, page 28-34](#)) also appear.

#### Related Topics

- [Sample URL, page 28-34](#)
- [Adding Basic HTTP Authentication credentials, page 28-34](#)
- [File Format, page 28-34](#)
- [Timestamps, page 28-34](#)
- [Keys, page 28-34](#)
- [Streaming, page 28-35](#)

### Sample URL

`http://example.com/monitor/content_filters?format=csv&sort_col_ss_0_0_0=MAIL_CONTENT_FILTER_INCOMING.RECIPIENTS_MATCHED&section=ss_0_0_0&date_range=current_day&sort_order_ss_0_0_0=desc&report_def_id=mga_content_filters`

### Adding Basic HTTP Authentication credentials

To specify basic HTTP Authentication credentials to the URL:

`http://example.com/monitor/`

becomes:

`http://username:password@example.com/monitor/`

### File Format

The downloaded file is in CSV format and has a .csv file extension. The file header has a default filename, which starts with the name of the report, then the section of the report.

### Timestamps

Exports that stream data show begin and end timestamps for each raw “interval” of time. Two begin and two end timestamps are provided — one in numeric format and the other in human-readable string format. The timestamps are in GMT time, which should make log aggregation easier if you have appliances in multiple time zones.

Note that in some rare cases where the data has been merged with data from other sources, the export file does not include timestamps. For example, the Outbreak Details export merges report data with Threat Operations Center (TOC) data, making timestamps irrelevant because there are no intervals.

### Keys

Exports also include the report table key(s), even in cases where the keys are not visible in the report. In cases where a key is shown, the display name shown in the report is used as the column header. Otherwise, a column header such as “key0,” “key1,” etc. is shown.

## Streaming

Most exports stream their data back to the client because the amount of data is potentially very large. However, some exports return the entire result set rather than streaming data. This is typically the case when report data is aggregated with non-report data (e.g. Outbreaks Detail.)

# Reporting Overview

Reporting in AsyncOS involves three basic actions:

- You can create Scheduled Reports to be run on a daily, weekly, or monthly basis.
- You can generate a report immediately (“on-demand” report).
- You can view archived versions of previously run reports (both scheduled and on-demand).

Configure scheduled and on-demand reports via the Monitor > Scheduled Reports page. View archived reports via the Monitor > Archived Reports page.

Your appliance will retain the most recent reports it generates, up to 1000 total versions for all reports. You can define as many recipients for reports as you want, including zero recipients. If you do not specify an email recipient, the system will still archive the reports. If you need to send the reports to a large number of addresses, however, it may be easier to create a mailing list rather than listing the recipients individually.

By default, the appliance archives the twelve most recent reports of each scheduled report. Reports are stored in the `/saved_reports` directory of the appliance. (See [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) for more information.)

## Related Topics

- [Scheduled Report Types, page 28-35](#)
- [Setting the Return Address for Reports, page 28-36](#)

## Scheduled Report Types

You can choose from the following report types:

- Content Filters
- Delivery Status
- DLP Incident Summary
- Executive Summary
- Incoming Mail Summary
- Internal Users Summary
- Outgoing Destinations
- Outgoing Mail Summary
- Outgoing Senders: Domains
- Sender Groups
- System Capacity
- TLS Connections

- Outbreak Filters
- Virus Types

Each of the reports consists of a summary of the corresponding Email Security Monitor page. So, for example, the Content Filters report provides a summary of the information displayed on the Monitor > Content Filters page. The Executive Summary report is based on the Monitor > Overview page.

**Related Topics**

- [Notes on Reports, page 28-36](#)

## Notes on Reports

Content Filter reports in a PDF format are limited to a maximum of 40 content filters. You can obtain the full listing via reports in a CSV format.

**Note**

To generate PDFs in Chinese, Japanese, or Korean on Windows computers, you must also download the applicable Font Pack from Adobe.com and install it on your local computer.

## Setting the Return Address for Reports

To set the return address for reports, see [Configuring the Return Address for Appliance Generated Messages, page 33-33](#). From the CLI, use the `addressconfig` command.

## Managing Reports

You can create, edit, delete, and view archived scheduled reports. You can also run a report immediately (on-demand report). The following report types are available: Content Filters, DLP Incident Summary, Executive Summary, Incoming Mail Summary, Internal Users Summary, Outgoing Mail Summary, Sender Groups, and Outbreak Filters. Managing and viewing these reports is discussed below.

**Note**

When in Cluster Mode, you are unable to view reports. You may view reports when in machine mode.

The Monitor > Scheduled Reports page shows a listing of the scheduled reports already created on the appliance.

**Related Topics**

- [Scheduled Reports, page 28-36](#)
- [Archived Reports, page 28-38](#)

## Scheduled Reports

Scheduled reports can be scheduled to run on a daily, weekly, or monthly basis. You can select a time at which to run the report. Regardless of when you run a report, it will only include data for the time period that you specify, for example the past 3 days or the previous calendar month. Note that a daily report scheduled to run at 1AM will contain data for the previous day, midnight to midnight.

Your appliance ships with a default set of scheduled reports—you can use, modify, or delete any of them.

#### Related Topics

- [Scheduling a Report to be Generated Automatically, page 28-37](#)
- [Editing Scheduled Reports, page 28-37](#)
- [Deleting Scheduled Reports, page 28-38](#)

## Scheduling a Report to be Generated Automatically

### Procedure

- 
- Step 1** On the Monitor > Scheduled Reports page, click **Add Scheduled Report**.
- Step 2** Select a report type. Depending on the report type you select, different options may be available. For more information about the available types of scheduled reports, see [Scheduled Report Types, page 28-35](#).
- Step 3** Enter a descriptive title for the report. AsyncOS does not verify the uniqueness of report names. To avoid confusion, do not create multiple reports with the same name.
- Step 4** Select a time range for the report data. (This option is not available for Outbreak Filters reports.)
- Step 5** Select a format for the report:
- **PDF.** Create a formatted PDF document for delivery, archival, or both. You can view the report as a PDF file immediately by clicking Preview PDF Report.  
For information about generating PDFs in languages other than English, see the [“Notes on Reports” section on page 28-36](#).
  - **CSV.** Create an ASCII text file that contains the tabular data as comma-separated values. Each CSV file may contain up to 100 rows. If a report contains more than one type of table, a separate CSV file is created for each table.
- Step 6** Specify the report options, if available. Some reports do not have report options.
- Step 7** Specify scheduling and delivery options. If you do not specify an email address, the report is archived but is not sent to any recipients.
-  **Note** If you are sending reports to an external account (such as Yahoo or Gmail, etc.), you may need to add the reporting return address to the external account’s whitelist to prevent report emails from being incorrectly classified as spam.
- 
- Step 8** Click **Submit**. Commit your changes.
- 

## Editing Scheduled Reports

### Procedure

- 
- Step 1** Click the report title in the listing on the Services > Centralized Reporting page.
- Step 2** Make your changes.



- Step 3** Submit and commit your changes.
- 

## Deleting Scheduled Reports

### Procedure

---

- Step 1** On the Services > Centralized Reporting page, select the check boxes corresponding to the reports that you want to delete.



**Note** Select the All check box to remove all scheduled reports.

---

- Step 2** Click **Delete**.
- Step 3** Confirm the deletion and then commit your changes.
- Any archived versions of deleted reports are *not* automatically deleted.
- 

## Archived Reports

The Monitor > Archived Reports page lists the available archived reports. You can view a report by clicking its name in the Report Title column. You can generate a report immediately by clicking **Generate Report Now**.

Use the Show menu to filter which type of reports is listed. Click the column headings to sort the listing.

Archived reports are deleted automatically — up to 30 instances of each scheduled report (up to 1000 reports) are kept and as new reports are added, older ones are deleted to keep the number at 1000. The 30 instances limit is applied to each individual scheduled report, not report type.

### Related Topics

- [Generating On-Demand Reports, page 28-38](#)

## Generating On-Demand Reports

You can generate a report without scheduling it. These on-demand reports are still based on a specified time frame, but they are generated immediately.

### Procedure

---

- Step 1** Click **Generate Report Now** on the Archived Reports page.
- Step 2** Select a report type and edit the title if desired. AsyncOS does not verify the uniqueness of report names. To avoid confusion, do not create multiple reports with the same name.
- For more information about the available types of scheduled reports, see [Scheduled Report Types, page 28-35](#).
- Step 3** Select a time range for the report data. (This option is not available for Virus Outbreak reports.)

If you create a custom range, the range will appear as a link. To modify the range, click the link.

**Step 4** Select a format for the report.

- **PDF.** Create a formatted PDF document for delivery, archival, or both. You can view the report as a PDF file immediately by clicking Preview PDF Report.

For information about generating PDFs in languages other than English, see the [“Notes on Reports” section on page 28-36](#).

- **CSV.** Create an ASCII text file that contains the tabular data as comma-separated values. Each CSV file may contain up to 100 rows. If a report contains more than one type of table, a separate CSV file is created for each table. Specify any report options.

**Step 5** Select whether to archive the report (if so, the report will shown on the Archived Reports page).

**Step 6** Specify whether to email the report and to which email addresses to send the report.

**Step 7** Click **Deliver this Report** to generate the report and deliver it to recipients or archive it.

**Step 8** Commit your changes.

---

## Troubleshooting Email Reports

**Problem** Drilling down from a report to view details in message tracking yields unexpected results.

**Solution** This can occur if reporting and message tracking were not simultaneously enabled, functioning properly, and storing data locally (as opposed to being stored centrally on a Security Management appliance). Data for each feature (reporting and message tracking) is stored only while that feature is enabled and functioning on that appliance, independently of whether the other feature (reporting or message tracking) is enabled and functioning. Therefore, reports may include data that is not available in Message Tracking and vice-versa.





## Tracking Messages

- [Message Tracking Overview, page 29-1](#)
- [Enabling Message Tracking, page 29-1](#)
- [Searching for Messages, page 29-2](#)
- [Working with Message Tracking Search Results, page 29-4](#)
- [Checking Message Tracking Data Availability, page 29-6](#)
- [Troubleshooting Message Tracking, page 29-7](#)

## Message Tracking Overview

Message tracking helps resolve help desk calls by giving a detailed view of message flow. For example, if a message was not delivered as expected, you can determine if it was found to contain a virus or placed in a spam quarantine — or if it is located somewhere else in the mail stream.

You can search for a particular email message or a group of messages that match criteria that you specify.



**Note**

You cannot use message tracking to read the content of messages.

## Enabling Message Tracking



**Note**

Message tracking data is preserved only for messages that are processed after you enable this feature.

### Before you Begin

- In order to search for and display attachment names in Message Tracking and view attachment names in log files, you must configure and enable at least one body scanning process, such as a message filter or content filter.
- To support searching by subject, log files must be configured to record subject headers. For more information, see [Chapter 38, “Logging.”](#)
- If you are setting up Centralized Tracking:  
Set up your Security Management appliance to support centralized message tracking for this Email Security appliance. See the *Cisco Content Security Management Appliance User Guide*.

### Procedure

- 
- Step 1** Click **Services > Centralized Services > Message Tracking**.  
Use this path even if you do not plan to centralize this service.
- Step 2** Select **Enable Message Tracking Service**.
- Step 3** If you are enabling message tracking for the first time after running the System Setup Wizard, review the end-user license agreement, and click **Accept**.
- Step 4** Choose a Message Tracking Service:

| Option               | Description                                                                                                      |
|----------------------|------------------------------------------------------------------------------------------------------------------|
| Local Tracking       | Use message tracking on this appliance.                                                                          |
| Centralized Tracking | Use a Security Management appliance to track messages for multiple Email Security appliances including this one. |

- Step 5** (Optional) Select the check box to save information for rejected connections.  
For best performance, leave this setting disabled.
- Step 6** Submit and commit your changes.
- 

### What To Do Next

If you selected Local Tracking:

- Choose who can access content related to DLP violations. See [Controlling Access to Sensitive Information in Message Tracking, page 32-5](#).
- (Optional) Adjust the disk space allocation for storing messages. See [Managing Disk Space, page 33-16](#).

## Searching for Messages

### Procedure

- 
- Step 1** Choose **Monitor > Message Tracking**.
- Step 2** Enter search criteria.
- To view all options, click the **Advanced** link.
  - Tracking does not support wildcard characters or regular expressions.
  - Tracking searches are not case sensitive.
  - Unless otherwise specified, the query is an “AND” search: The query returns messages that match *all* conditions specified in the search fields. For example, if you specify text strings for the envelope recipient and the subject line parameters, the query returns only messages that match *both* the specified envelope recipient *and* the subject line.
  - Search criteria include:

| Option                                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Envelope Sender</b>                           | <p>Select <b>Begins With</b>, <b>Is</b>, or <b>Contains</b>, then enter an email address, username, or domain of a message sender to find.</p> <p>You can enter any character(s). No validation of your entry is performed.</p>                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Envelope Recipient</b>                        | <p>Select <b>Begins With</b>, <b>Is</b>, or <b>Contains</b>, and enter an email address, username, or domain of a message recipient to find.</p> <p>You can enter any character(s). No validation of your entry is performed.</p>                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Subject</b>                                   | <p>Select <b>Begins With</b>, <b>Is</b>, or <b>Contains</b>, and enter a text string to search for in the message subject line.</p> <p><b>Warning:</b> Do not use this type of search in environments where regulations prohibit such tracking.</p>                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Message Received</b>                          | <p>Specify a date and time range.</p> <p>If you do not specify a date, the query returns data for all dates. If you specify a time range only, the query returns data for that time range across all available dates.</p> <p>Use the local date and time that the message was received by the Email Security appliance.</p>                                                                                                                                                                                                                                                                                                                  |
| <b>Advanced options:</b>                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Sender IP Address/ Domain / Network Owner</b> | <p>Specify the IP address, domain, or network owner of a remote host.</p> <p>You can search within rejected connections only or search all messages.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Attachment</b>                                | <p>Select <b>Begins With</b>, <b>Is</b>, or <b>Contains</b>, and enter an ASCII or Unicode text string for one attachment to find. Leading and trailing spaces are <i>not</i> stripped from the text you enter.</p> <p>You can search for messages by attachment filenames only if you have performed:</p> <ul style="list-style-type: none"> <li>• Body scan using a message filter</li> <li>• Body scan using a content filter</li> <li>• Advanced Malware Protection (AMP) scan.</li> </ul> <p>For more information about identifying files based on SHA-256 hash, see <a href="#">Identifying Files by SHA-256 Hash, page 16-11</a>.</p> |
| <b>Message Event</b>                             | <p>Select one or more message processing events. For example, you can search for messages that have been delivered, quarantined, or hard bounced.</p> <p>Message events are added with an “OR” operator: Selecting multiple events finds messages that match <i>any</i> of the conditions you specify.</p>                                                                                                                                                                                                                                                                                                                                   |
| <b>Message ID Header</b>                         | <p>Enter a text string for the SMTP Message-ID header.</p> <p>This RFC 822 message header uniquely identifies each email message. It is inserted in the message when the message is first created.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Cisco IronPort MID</b>                        | <p>Enter a message number to search for. An IronPort MID uniquely identifies each email message on the Email Security appliance.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Query Settings</b>                            | <p>Change the default query timeout and maximum number of results to return.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

- Step 3** Click **Search** to submit the query.
- The query results are displayed at the bottom of the page.
- 

#### Related Topics

- [Working with Message Tracking Search Results, page 29-4](#)

## Working with Message Tracking Search Results

Keep the following points in mind:

- Search results depend on your configuration. For example, if you search for messages in a URL Category for which you have not filtered, you will find no results.
- For information about searches involving Advanced Malware Protection (file reputation scanning and file analysis), see [About Message Tracking and Advanced Malware Protection Features, page 16-12](#).

Actions you can take when working with search results:

- Show more than 250 search results by returning to the search criteria, clicking Advanced, scrolling to the Query Settings, and setting the maximum number of results to 1000.
- Show more results per page by choosing an option from the top right side of the search results section.
- Navigate through multiple pages of search results from the top right side of the search results section.
- Narrow your search results by floating the cursor over a value in the search results that you want to add as a condition. If an orange highlight appears, you can click that value to narrow the search by that criterion. This adds the additional criterion to the search criteria. For example, if you search for messages sent to a particular recipient, you can then click on a sender name in the search results to find all messages to that recipient from that sender within the time range (and meeting any other criteria) that you originally specified.
- If more than 1000 messages match your search criteria, you can click Export All (a link at the top right of the search results section) and export up to 50,000 search results as a comma-separated values file and work with the data in another application.
- View more details for a message by clicking Show Details in the row for that message. A new browser window opens with the message details.
- For quarantined messages, you can click a link in the message tracking search results to view details such as the reason the message was quarantined.



#### Note

If you clicked a link in a report page to view message details in Message Tracking, and the set of results is not what you expected, this can occur if reporting and tracking were not both simultaneously and continuously enabled during the time period you are reviewing.

---

#### Related Topics

- [Message Details, page 29-5](#)

## Message Details

| Item                                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Envelope and Header Summary section:</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Received Time</b>                        | Time that the Email Security appliance received the message.<br><br>Dates and times are displayed using the local time configured on the Email Security appliance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>MID</b>                                  | Unique IronPort message ID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Message Size</b>                         | Message size.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Subject</b>                              | Subject line of the message.<br><br>The subject line in the tracking results may have the value “(No Subject)” if the message does not have a subject, or if log files are not configured to record subject headers. For more information, see <a href="#">Chapter 38, “Logging.”</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Envelope Sender</b>                      | Address of the sender in the SMTP envelope.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Envelope Recipients</b>                  | If your deployment uses the alias table for alias expansion, the search finds the expanded recipient addresses rather than the original envelope addresses. For more information about Alias Tables, see “Creating Alias Tables” in the “Configuring Routing and Delivery Features” chapter .<br><br>In all other cases, message tracking queries find the original envelope recipient addresses.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Message ID Header</b>                    | The RFC 822 message header.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>SMTP Auth User ID</b>                    | SMTP authenticated username of the sender, if the sender used SMTP authentication to send the message. Otherwise, the value is “N/A.”                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Attachments</b>                          | The names of files attached to the message.<br><br>Messages that contain at least one attachment with the queried name will appear in the search results.<br><br>Some attachments may not be tracked. For performance reasons, scanning of attachment names occurs only as part of other scanning operations, for example message or content filtering, DLP, or disclaimer stamping. Attachment names are available only for messages that pass through body scanning while the attachment is still attached. Situations in which an attachment name will not appear in search results include (but are not limited to): <ul style="list-style-type: none"> <li>• if the system only uses content filters, and a message is dropped or its attachment is stripped by anti-spam or anti-virus filters</li> <li>• if message splintering policies strip the attachment from some messages before body scanning occurs.</li> </ul> For performance reasons, the names of files within attachments, such as OLE objects or archives such as .ZIP files, are not searched. |
| <b>Sending Host Summary section</b>         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Reverse DNS Hostname</b>                 | Name of the sending host, as verified by reverse DNS (PTR) lookup.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>IP Address</b>                           | IP address of the sending host.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |



| Item                                                                                                                                             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SBRS Score</b>                                                                                                                                | <p>SenderBase reputation score. The range is from 10 (likely a trustworthy sender) to -10 (apparent spammer). A score of “None” indicates that there was no information about this host at the time the message was processed.</p> <p>For more information about SBRS, see <a href="#">Chapter 6, “Sender Reputation Filtering.”</a></p>                                                                                                                                         |
| <b>Processing Details section</b>                                                                                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Summary information</b><br><br>(A Summary tab displays only if there is also a DLP Matched Content tab. Summary information always displays.) | <p>The Summary section displays status events logged during the processing of the message.</p> <p>Entries include information about Mail Policy processing, such as Anti-Spam and Anti-Virus scanning, and other events such as message splitting and custom log entries added by a content or message filter.</p> <p>If the message was delivered, the details of the delivery are displayed here.</p> <p>The last recorded event is highlighted in the processing details.</p> |
| <b>DLP Matched Content tab</b>                                                                                                                   | <p>This tab displays only for messages that were caught by DLP policies.</p> <p>It includes information about the match, as well as the sensitive content that triggered the DLP policy match.</p> <p>You can control who has access to this information. See <a href="#">Controlling Access to Sensitive Information in Message Tracking, page 32-5</a>.</p>                                                                                                                    |

**Related Topics**

- [Searching for Messages, page 29-2](#)

## Checking Message Tracking Data Availability

You can determine the date range that your message tracking data includes, as well as identify any missing intervals in that data.

**Procedure**

- 
- Step 1** Select **Monitor > Message Tracking**.
- Step 2** Look for **Data in time range:** in the upper right corner of the Search box.
- Step 3** Click the value shown for **Data in time range:**.
- 

**Related Topics**

- [About Message Tracking and Upgrades, page 29-6](#)

## About Message Tracking and Upgrades

New message tracking features may not apply to messages that were processed before upgrade, because the required data may not have been retained for those messages. For possible limitations related to message tracking data and upgrades, see the Release Notes for your release.

# Troubleshooting Message Tracking

## Related Topics

- [Attachments Do Not Appear in Search Results, page 29-7](#)
- [Expected Messages Are Missing from Search Results, page 29-7](#)

## Attachments Do Not Appear in Search Results

**Problem** Attachment names are not found and displayed in search results.

**Solution** See configuration requirements at [Enabling Message Tracking, page 29-1](#). Also see limitations for attachment name searches in [Message Details, page 29-5](#).

## Expected Messages Are Missing from Search Results

**Problem** Search results did not include messages that should have met the criteria.

### Solution

- Results for many searches, and especially searches that involve Message Events, depend on your appliance configuration. For example, if you search for a URL Category for which you have not filtered, no results will be found, even if a message contains a URL in that category. Verify that you have configured the Email Security appliance properly to achieve the behavior that you expected. For example, check your mail policies, content and message filters, and quarantine settings.
- If expected information is missing after you clicked a link in a report, see [Troubleshooting Email Reports, page 28-39](#).





## Policy, Virus, and Outbreak Quarantines

---

- [Overview of Policy, Virus, and Outbreak Quarantines, page 30-1](#)
- [Managing Policy, Virus, and Outbreak Quarantines, page 30-3](#)
- [Working with Messages in Policy, Virus, or Outbreak Quarantines, page 30-10](#)

### Overview of Policy, Virus, and Outbreak Quarantines

“Policy, virus and outbreak quarantines” includes all non-spam quarantines, including the File Analysis quarantine.

When an Email Security appliance detects possible malware or content that is not allowed by your organization in incoming or outgoing messages, it can send those messages to a quarantine instead of deleting them immediately. A quarantine holds these messages safely on the Email Security appliance or on a Cisco Content Security Management appliance for a period of time, to allow a human being to review them, or to await an update that will better evaluate the safety of the message.

Examples of how non-spam quarantines can be used in your organization:

- **Policy enforcement.** Let Human Resources personnel or the Legal department review messages that may contain offensive, confidential, or otherwise disallowed information.
- **Virus quarantine.** Store messages that are marked as infected, encrypted, or not scannable by the anti-virus scanning engine to prevent the spread of viruses to your users.
- **Outbreak prevention.** Hold messages that are flagged by the Outbreak Filters as possibly being part of a viral outbreak or small-scale malware attack until an anti-virus or anti-spam update is released.
- **File Analysis quarantine.** Store messages that have attachments that may contain malware, and that have been sent for analysis, until a verdict is reached.

#### Related Topic

- [Quarantine Types, page 30-2](#)
- [Chapter 31, “Spam Quarantine”](#)

## Quarantine Types

| Quarantine Type             | Quarantine Name                      | Created by the System by Default? | Description                                                                                                                                                                                                                                                          | More Information                                                                                                                                                                                                                       |
|-----------------------------|--------------------------------------|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Advanced Malware Protection | File Analysis                        | Yes                               | <p>Holds messages that are sent for file analysis, until a verdict is returned.</p> <p>For special characteristics, see:</p> <ul style="list-style-type: none"> <li>• <a href="#">Quarantining Messages with Attachments Sent for Analysis, page 16-7</a></li> </ul> | <ul style="list-style-type: none"> <li>• <a href="#">Managing Policy, Virus, and Outbreak Quarantines, page 30-3</a></li> <li>• <a href="#">Working with Messages in Policy, Virus, or Outbreak Quarantines, page 30-10</a></li> </ul> |
| Virus                       | Virus                                | Yes                               | Holds messages that may be transmitting malware, as determined by the anti-virus engine.                                                                                                                                                                             |                                                                                                                                                                                                                                        |
| Outbreak                    | Outbreak                             | Yes                               | Holds messages caught by Outbreak Filters as potentially being spam or malware.                                                                                                                                                                                      |                                                                                                                                                                                                                                        |
| Policy                      | Policy                               | Yes                               | <p>Holds messages caught by message filters, content filters, and DLP message actions.</p> <p>A default Policy quarantine has been created for you.</p>                                                                                                              |                                                                                                                                                                                                                                        |
|                             | Unclassified                         | Yes                               | <p>Holds messages only if a quarantine that is specified in a message filter, content filter, or DLP message action has been deleted.</p> <p>You cannot assign this quarantine to any filter or message action.</p>                                                  |                                                                                                                                                                                                                                        |
|                             | (Policy quarantines that you create) | No                                | Policy quarantines that you create for use in message filters, content filters, and DLP message actions.                                                                                                                                                             |                                                                                                                                                                                                                                        |
| Spam                        | Spam                                 | Yes                               | <p>Holds spam or suspected spam messages for the message's recipient or an administrator to review.</p> <p>The spam quarantine is not included in the group of policy, virus, and outbreak quarantines and is managed separately from all other quarantines.</p>     | <a href="#">Chapter 31, "Spam Quarantine"</a>                                                                                                                                                                                          |

# Managing Policy, Virus, and Outbreak Quarantines

- [Disk Space Allocation for Policy, Virus, and Outbreak Quarantines, page 30-3](#)
- [Retention Time for Messages in Quarantines, page 30-3](#)
- [Default Actions for Automatically Processed Quarantined Messages, page 30-4](#)
- [Checking the Settings of System-Created Quarantines, page 30-5](#)
- [Configuring Policy, Virus, and Outbreak Quarantines, page 30-5](#)
- [About Editing Policy, Virus, and Outbreak Quarantine Settings, page 30-6](#)
- [Determining the Filters and Message Actions to Which a Policy Quarantine Is Assigned, page 30-7](#)
- [About Deleting Policy Quarantines, page 30-7](#)
- [Monitoring Quarantine Status, Capacity, and Activity, page 30-8](#)
- [Policy Quarantine Performance, page 30-8](#)
- [Alerts About Quarantine Disk-Space Usage, page 30-9](#)
- [Policy Quarantines and Logging, page 30-9](#)
- [About Distributing Message Processing Tasks to Other Users, page 30-9](#)
- [About Policy, Virus, and Outbreak Quarantines in Cluster Configurations, page 30-10](#)
- [About Centralized Policy, Virus, and Outbreak Quarantines, page 30-10](#)

## Disk Space Allocation for Policy, Virus, and Outbreak Quarantines

For disk space information for policy, virus, and outbreak quarantines, see [Managing Disk Space, page 33-16](#).

Policy, virus, and outbreak quarantines consume some disk space on the Email Security appliance even if the quarantines are centralized.

Messages in multiple quarantines consume the same amount of disk space as a message in a single quarantine.

### Related Topics

- [Monitoring Quarantine Status, Capacity, and Activity, page 30-8](#)
- [Alerts About Quarantine Disk-Space Usage, page 30-9](#)
- [Retention Time for Messages in Quarantines, page 30-3](#)

## Retention Time for Messages in Quarantines

Messages are automatically removed from the quarantine under the following circumstances:

- **Normal Expiration**—the configured retention time is met for a message in the quarantine. You specify a retention time for messages in each quarantine. Each message has its own specific expiration time, displayed in the quarantine listing. Messages are stored for the amount of time specified unless another circumstance described in this topic occurs.

**Note**

The normal retention time for messages in the Outbreak Filters quarantine is configured in the Outbreak Filters section of each mail policy, not in the outbreak quarantine. For information, see [Chapter 14, “Outbreak Filters.”](#)

- Early Expiration—messages are forced from quarantines before the configured retention time is reached. This can happen when:

- The size limit for all quarantines, as defined in [Disk Space Allocation for Policy, Virus, and Outbreak Quarantines, page 30-3](#), is reached.

If the size limit is reached, the oldest messages, regardless of quarantine, are processed and the default action is performed for each message, until the size of all quarantines is again less than the size limit. The policy is First In First Out (FIFO). Messages in multiple quarantines will be expired based on their latest expiration time.

(Optional) You can configure individual quarantines to be exempt from release or deletion because of insufficient disk space. If you configure all quarantines to be exempt and the disk space reaches capacity, messages in the quarantine will be delivered to make room for new messages.

You will receive alerts at disk-space milestones. See [Alerts About Quarantine Disk-Space Usage, page 30-9](#).

- You delete a quarantine that still holds messages.

When a message is automatically removed from a quarantine, the default action is performed on that message. See [Default Actions for Automatically Processed Quarantined Messages, page 30-4](#).

**Note**

In addition to the above scenarios, messages can be automatically removed from quarantine based on the result of scanning operations (outbreak filters or file analysis.)

**Effects of Time Adjustments on Retention Time**

- Daylight savings time and appliance time zone changes do not affect the retention period.
- If you change the retention time of a quarantine, only new messages will have the new expiration time.
- If the system clock is changed, messages that should have expired in the past will expire at the next most appropriate time.
- System clock changes do not apply to messages that are in the process of being expired.

## Default Actions for Automatically Processed Quarantined Messages

The default action is performed on messages in a policy, virus, or outbreak quarantine when any situation described in [Retention Time for Messages in Quarantines, page 30-3](#), occurs.

There are two primary default actions:

- Delete—The message is deleted.
- Release—The message is released for delivery.

Upon release, messages may be rescanned for threats. For more information, see [About Rescanning of Quarantined Messages, page 30-17](#).

In addition, messages released before their expected retention time has passed can have additional operations performed on them, such as adding an X-Header. For more information, see [Configuring Policy, Virus, and Outbreak Quarantines, page 30-5](#).

## Checking the Settings of System-Created Quarantines

Before you use quarantines, customize the settings of the default quarantines, including the Unclassified quarantine.

### Related Topics

- [Configuring Policy, Virus, and Outbreak Quarantines, page 30-5](#)

## Configuring Policy, Virus, and Outbreak Quarantines

### Before You Begin

- If you are editing an existing quarantine, see [About Editing Policy, Virus, and Outbreak Quarantine Settings, page 30-6](#).
- Understand how messages in quarantines are automatically managed, including retention times and default actions. See [Retention Time for Messages in Quarantines, page 30-3](#), and [Default Actions for Automatically Processed Quarantined Messages, page 30-4](#).
- Determine which users you want to have access to each quarantine, and create users and custom user roles accordingly. For details, see [Which User Groups Can Access Policy, Virus, and Outbreak Quarantines, page 30-10](#).

### Procedure

---

**Step 1** Choose **Monitor > Policy, Virus, and Outbreak Quarantines**.

**Step 2** Do one of the following:

- Click **Add Policy Quarantine**.
- Click a quarantine to edit.

**Step 3** Enter information.

Keep the following in mind:

- If you do *not* want messages in this quarantine to be processed before the end of the Retention Period you specify, even when quarantine disk space is full, deselect **Free up space by applying default action on messages upon space overflow**.  
  
Do not select this option for all quarantines. The system must be able to make space by deleting messages from at least one quarantine.
- If you select **Release** as the default action, you can specify additional actions to apply to messages that are released before their retention period has passed:



| Option            | Information                                                                                                                                                                                                                                                                                                                                                                |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Modify Subject    | <p>Type the text to add and specify whether to add it to the beginning or the end of the original message subject.</p> <p>For example, you might want to warn the recipient that the message may contain inappropriate content.</p> <p><b>Note</b> In order for a subject with non-ASCII characters to display correctly it must be represented according to RFC 2047.</p> |
| Add X-Header      | <p>An X-Header can provide a record of actions taken on a message. This can be helpful for example when handling inquiries about why a particular message was delivered.</p> <p>Enter a name and value.</p> <p>Example:</p> <p>Name =Inappropriate-release-early</p> <p>Value = True</p>                                                                                   |
| Strip Attachments | Stripping attachments protects against viruses that may be in such files.                                                                                                                                                                                                                                                                                                  |

**Step 4** Specify the users who can access this quarantine:

| User                           | Information                                                                                                                                                                                                          |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local Users                    | <p>The list of local users includes only users with roles that can access quarantines.</p> <p>The list excludes users with Administrator privileges, because all Administrators have full access to quarantines.</p> |
| Externally Authenticated Users | You must have configured external authentication.                                                                                                                                                                    |
| Custom User Roles              | You see this option only if you have created at least one custom user role with quarantine access.                                                                                                                   |

**Step 5** Submit and commit your changes.

### What To Do Next

Create message and content filters and DLP message actions that will move messages to the quarantine. See [Chapter 9, “Using Message Filters to Enforce Email Policies”](#), [Chapter 11, “Content Filters”](#), and [Message Actions, page 17-34](#).

## About Editing Policy, Virus, and Outbreak Quarantine Settings



### Note

- You cannot rename a quarantine.

- See also [Effects of Time Adjustments on Retention Time](#), page 30-4.

To change quarantine settings, choose Monitor > Policy, Virus, and Outbreak Quarantines, and then click the name of a quarantine.

## Determining the Filters and Message Actions to Which a Policy Quarantine Is Assigned

You can view the message filters, content filters, Data Loss Prevention (DLP) message actions, and DMARC verification profiles that are associated with a policy quarantine.

### Procedure

- 
- |               |                                                                                                                       |
|---------------|-----------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Click <b>Monitor &gt; Policy, Virus, and Outbreak Quarantines</b> .                                                   |
| <b>Step 2</b> | Click the name of the policy quarantine to check.                                                                     |
| <b>Step 3</b> | Scroll to the bottom of the page and view the <b>Associated Message Filters/Content Filters/DLP Message Actions</b> . |
- 

## About Deleting Policy Quarantines

- Before you delete a policy quarantine, see if it is associated with any active filters or message actions. See [Determining the Filters and Message Actions to Which a Policy Quarantine Is Assigned](#), page 30-7.
- You can delete a policy quarantine even if it is assigned to a filter or message action.
- If you delete a quarantine that is not empty, the default action defined in the quarantine will be applied to all messages, even if you have selected the option not to delete messages if the disk is full. See [Default Actions for Automatically Processed Quarantined Messages](#), page 30-4.
- After you delete the quarantine associated with a filter or message action, any messages subsequently quarantined by that filter or message action will be sent to the Unclassified quarantine. You should customize the default settings of the Unclassified quarantine before you delete quarantines.
- You cannot delete the Unclassified quarantine.

## Monitoring Quarantine Status, Capacity, and Activity

| To View                                                                                                  | Do This                                                                                                                                                                                                                                                                    |
|----------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Currently available space for all non-spam quarantines                                                   | Choose <b>Monitor &gt; Policy, Virus, and Outbreak Quarantines</b> and look just below the table.                                                                                                                                                                          |
| Total amount of space currently used by all quarantines                                                  | Choose <b>Monitor &gt; System Status</b> and look for <b>Queue Space Used by Quarantine</b> .                                                                                                                                                                              |
| Amount of space currently used by each quarantine                                                        | Choose <b>Monitor &gt; Policy, Virus, and Outbreak Quarantines</b> , click the quarantine name, and look for this information in the table row directly below the quarantine name.                                                                                         |
| Total number of messages currently in all quarantines                                                    | Choose <b>Monitor &gt; System Status</b> and look for <b>Active Messages in Quarantine</b> .                                                                                                                                                                               |
| Number of messages currently in each quarantine                                                          | Choose <b>Monitor &gt; Policy, Virus, and Outbreak Quarantines</b> and look at the table row for the quarantine.                                                                                                                                                           |
| Total CPU usage by all quarantines                                                                       | Choose <b>Monitor &gt; System Status</b> and look in the <b>CPU Utilization</b> section.                                                                                                                                                                                   |
| Date and time when the last message entered each quarantine (excluding moves between policy quarantines) | Choose <b>Monitor &gt; Policy, Virus, and Outbreak Quarantines</b> and look at the table row for the quarantine.                                                                                                                                                           |
| Date a policy quarantine was created                                                                     | Choose <b>Monitor &gt; Policy, Virus, and Outbreak Quarantines</b> , click the quarantine name, and look for this information in the table row directly below the quarantine name.<br><br>Creation date and creator name are not available for system-created quarantines. |
| Name of policy quarantine creator                                                                        |                                                                                                                                                                                                                                                                            |
| Filters and message actions associated with a policy quarantine                                          | See <a href="#">Determining the Filters and Message Actions to Which a Policy Quarantine Is Assigned</a> , page 30-7.                                                                                                                                                      |

## Policy Quarantine Performance

Messages stored in policy quarantines use system memory in addition to hard-drive space. Storing hundreds of thousands of messages in policy quarantines on a single appliance may cause a decrease in the appliance's performance due to excessive memory usage. The appliance takes more time to quarantine, delete, and release messages, which causes message processing to slow down and the email pipeline to back up.

Cisco recommends storing an average of less than 20,000 messages in your policy quarantines to ensure that the Email Security appliance processes email at a normal rate.

To check the number of messages in quarantines, see [Monitoring Quarantine Status, Capacity, and Activity](#), page 30-8.

## Alerts About Quarantine Disk-Space Usage

An alert is sent whenever the total size of the policy, virus, and outbreak quarantine reaches or passes 75 percent, 85 percent, and 95 percent of its capacity. The check is performed when a message is placed in the quarantine. For example, if adding a message to a quarantine increases the size to or past 75 percent of the total capacity, an alert is sent. See also [Alerts, SNMP Traps, and Outbreak Filters](#), page 14-23.

For more information about Alerts, see [Alerts](#), page 33-34.

## Policy Quarantines and Logging

AsyncOS individually logs all messages that are quarantined:

```
Info: MID 482 quarantined to "Policy" (message filter:policy_violation)
```

The message filter or Outbreak Filters feature rule that caused the message to be quarantined is placed in parentheses. A separate log entry is generated for each quarantine in which the message is placed.

AsyncOS also individually logs messages that are removed from quarantine:

```
Info: MID 483 released from quarantine "Policy" (queue full)
```

```
Info: MID 484 deleted from quarantine "Anti-Virus" (expired)
```

The system individually logs messages after they are removed from all quarantines and either permanently deleted or scheduled for delivery, for example

```
Info: MID 483 released from all quarantines
```

```
Info: MID 484 deleted from all quarantines
```

When a message is re-injected, the system creates a new Message object with a new Message ID (MID). This is logged using an existing log message with a new MID “byline”, for example:

```
Info: MID 483 rewritten to 513 by Policy Quarantine
```

## About Distributing Message Processing Tasks to Other Users

You can distribute message review and processing tasks to other administrative users. For example:

- The Human Resources team can review and manage the Policy Quarantine.
- The Legal team can manage the Confidential Material Quarantine.

You assign access privileges to these users when you specify settings for a quarantine. In order to add users to quarantines, the users must already exist.

Each user may have access to all, some, or none of the quarantines. A user who is not authorized to view a quarantine will not see any indication of its existence anywhere in the GUI or CLI listings of quarantines.

### Related Topics

- [Which User Groups Can Access Policy, Virus, and Outbreak Quarantines](#), page 30-10
- [Working with User Accounts](#), page 32-1
- [External Authentication](#), page 32-20
- [Managing Custom User Roles for Delegated Administration](#), page 32-7

## Which User Groups Can Access Policy, Virus, and Outbreak Quarantines

When you allow administrative users to access a quarantine, the actions that they can perform depend on their user group:

- Users in the Administrators group can create, configure, delete, and centralize quarantines and can manage quarantined messages.
- Users in the Operators, Guests, Read-Only Operators, and Help Desk Users groups, as well as custom user roles with quarantine management privileges, can search for, view, and process messages in a quarantine, but cannot change the quarantine's settings, create, delete, or centralize quarantines. You specify in each quarantine which of these users have access to that quarantine.
- Users in the Technicians group cannot access quarantines.

Access privileges for related features, such as Message Tracking and Data Loss Prevention, also affect the options and information that an administrative user sees on Quarantine pages. For example, if a user does not have access to Message Tracking, that user will not see message tracking links and information for quarantined messages.

End users do not have see or have access to policy, virus, and outbreak quarantines.

## About Policy, Virus, and Outbreak Quarantines in Cluster Configurations

Policy, virus, and outbreak quarantines are configurable only at machine level in deployments with centralized management.

## About Centralized Policy, Virus, and Outbreak Quarantines

You can centralize policy, virus, and outbreak quarantines on a Cisco Content Security Management appliance. For information, see [About Centralizing Policy, Virus, and Outbreak Quarantines, page 42-5](#) and the user documentation for your Security Management appliance.

# Working with Messages in Policy, Virus, or Outbreak Quarantines

### Related Topics

- [Viewing Messages in Quarantines, page 30-11](#)
- [Finding Messages in Policy, Virus, and Outbreak Quarantines, page 30-11](#)
- [Manually Processing Messages in a Quarantine, page 30-12](#)
- [Messages in Multiple Quarantines, page 30-13](#)
- [Message Details and Viewing Message Content, page 30-14](#)
- [About Rescanning of Quarantined Messages, page 30-17](#)
- [The Outbreak Quarantine, page 30-17](#)

## Viewing Messages in Quarantines

| To                                                         | Do This                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| View all messages in a quarantine                          | <p>Choose <b>Monitor &gt; Policy, Virus, and Outbreak Quarantines</b>.</p> <p>In the row for the relevant quarantine, click the blue number in the <b>Messages</b> column of the table.</p>                                                                                                                                                           |
| View messages in the Outbreak quarantine                   | <ul style="list-style-type: none"> <li>Choose <b>Monitor &gt; Policy, Virus, and Outbreak Quarantines</b>.</li> </ul> <p>In the row for the relevant quarantine, click the blue number in the <b>Messages</b> column of the table.</p> <ul style="list-style-type: none"> <li>See <a href="#">Manage by Rule Summary Link, page 30-18</a>.</li> </ul> |
| Navigate through the list of messages in a quarantine      | Click Previous, Next, a page number, or double-arrow link. The double arrows take you to the first (<<) or last (>>) page in the listing.                                                                                                                                                                                                             |
| Sort the list of messages in a quarantine                  | Click a column heading (except columns that could include multiple items or the “In other quarantines” column).                                                                                                                                                                                                                                       |
| Resize table columns                                       | Drag the divider between column headings.                                                                                                                                                                                                                                                                                                             |
| View the content that caused the message to be quarantined | See <a href="#">Viewing Matched Content, page 30-15</a> .                                                                                                                                                                                                                                                                                             |

### Related Topics

- [Quarantined Messages and International Character Sets, page 30-11](#)

## Quarantined Messages and International Character Sets

For messages with subjects that contain characters from international character sets (double-byte, variable length, and non-ASCII encoded), the Policy Quarantine pages display subject lines in non-ASCII characters in their decoded form.

## Finding Messages in Policy, Virus, and Outbreak Quarantines



### Note

- Searches in Policy, Virus, and Outbreak quarantines do not find messages in the spam quarantine.
- Users can find and see only the messages in quarantines to which they have access.

### Procedure

- Step 1** Choose **Monitor > Policy, Virus, and Outbreak Quarantines**.
- Step 2** Click the **Search Across Quarantines** button.

**Tip**

For the Outbreak Quarantine, you can also find all messages quarantined by each outbreak rule: Click the **Manage by Rule Summary** link in the Outbreak table row, and then click the relevant rule.

**Step 3** Select the quarantines in which to search.

**Step 4** (Optional) Enter other search criteria.

- For Envelope Sender and Envelope Recipient: You can enter any character(s). No validation of your entry is performed.
- Search results include only messages that match *all* of the criteria you specify. For example, if you specify an Envelope Recipient and a Subject, only messages that match the terms specified in both the Envelope Recipient *and* the Subject are returned.

**What To Do Next**

You can use the search results in the same way that you use the quarantine listings. For more information, see [Manually Processing Messages in a Quarantine](#), page 30-12.

## Manually Processing Messages in a Quarantine

Manually processing messages means to manually select a Message Action for the message from the Message Actions page.

**Note**

For deployments with RSA Enterprise Manager, you can view quarantined messages on the Email Security appliance or on Enterprise Manager, but you must use Enterprise Manager to take action on messages. For information about Enterprise Manager, see [Chapter 17, “Data Loss Prevention.”](#)

You can perform the following actions on messages:

- Delete
- Release
- Delay Scheduled Exit from quarantine
- Send a Copy of messages to email addresses that you specify
- Move a message from one quarantine to another

Generally, you can perform actions on messages in the lists that are displayed when you do the following. However, not all actions are available in all situations.

- From the list of quarantines on the Monitor > Policy, Virus, and Outbreak Quarantines page, click the number of messages in a quarantine.
- Click Search Across Quarantines.
- Click a quarantine name and search within a quarantine.

You can perform these actions on multiple messages at one time by:

- Choosing an option from the pick list at the top of the list of messages.
- Selecting the check box beside each message listed on a page.
- Selecting the check box in the table heading at the top of a list of messages. This applies the action to all messages visible on the screen. Messages on other pages are not affected.

Additional options are available for messages in the outbreak quarantine. See [Outbreak Quarantine and the Manage by Rule Summary View](#), page 14-22.

#### Related Topics

- [Sending a Copy of the Message](#), page 30-13
- [About Moving Messages Between Policy Quarantines](#), page 30-13
- [Messages in Multiple Quarantines](#), page 30-13
- [Default Actions for Automatically Processed Quarantined Messages](#), page 30-4

## Sending a Copy of the Message

Only users who belong to the Administrators group may send copies of a message.

To send a copy of the message, enter an email address in the Send Copy To: field and click **Submit**. Sending a copy of a message does not cause any other action to be performed on the message.

## About Moving Messages Between Policy Quarantines

You can manually move messages from one policy quarantine to another on a single appliance.

When you move a message to a different quarantine:

- The expiration time is unchanged. The message keeps the retention time of the original quarantine.
- The reason the message was quarantined, including the matched content and other relevant details, does not change.
- If a message is in multiple quarantines and you move the message to a destination that already holds a copy of that message, the expiration time and reason for quarantine of the moved copy of the message overwrite those of the copy of the message that was originally in the destination quarantine.

## Messages in Multiple Quarantines

If a message is present in one or more other quarantines, the “In other quarantines” column in the quarantine message list will show “Yes,” regardless of whether you have permissions to access those other quarantines.

A message in multiple quarantines:

- Is not delivered unless it has been released from all of the quarantines in which it resides. If it is deleted from any quarantine, it will never be delivered.
- Is not deleted from any quarantine until it has been deleted or released from all quarantines in which it resides.

Because a user wanting to release a message may not have access to all of the quarantines in which it resides, the following rules apply:



- A message is not released from any quarantine until it has been released from all of the quarantines in which it resides.
- If a message is marked as Deleted in any quarantine, it cannot be delivered from any other quarantine in which it resides. (It can still be released.)

If a message is queued in multiple quarantines and a user does not have access to one or more of the other quarantines:

- The user will be informed whether the message is present in each of the quarantines to which the user has access.
- The GUI shows only the scheduled exit time from the quarantines to which the user has access. (For a given message, there is a separate exit time for each quarantine.)
- The user will not be told the names of the other quarantine(s) holding the message.
- The user will not see matched content that caused the message to be placed into quarantines that the user does not have access to.
- Releasing a message affects only the queues to which the user has access.
- If the message is also queued in other quarantines not accessible to the user, the message will remain in quarantine, unchanged, until acted upon by users who have the required access to the remaining quarantines (or until the message is released “normally” via early or normal expiration).

## Message Details and Viewing Message Content

Click on the subject line of a message to view that message’s content and to access the Quarantined Message page.

The Quarantined Message page has two sections: Quarantine Details and Message Details.

From the Quarantined Message page, you can read the message, select a Message Action, send a copy of the message, or test for viruses. You can also see if a message will be encrypted upon release from the quarantine due to the Encrypt on Delivery filter action.

The Message Details section displays the message body, message headers, and attachments. Only the first 100 K of the message body is displayed. If the message is longer, the first 100 K is shown, followed by an ellipsis (...). The actual message is not truncated. This is for display purposes only. You can download the message body by clicking [message body] in the Message Parts section at the bottom of Message Details. You can also download any of the message’s attachments by clicking the attachment’s filename.

If you view a message that contains a virus and you have desktop anti-virus software installed on your computer, your anti-virus software may complain that it has found a virus. This is not a threat to your computer and can be safely ignored.

To view additional details about the message, click the **Message Tracking** link.



### Note

For the special Outbreak quarantine, additional functionality is available. See [The Outbreak Quarantine, page 30-17](#).

### Related Topics

- [Viewing Matched Content, page 30-15](#)
- [Downloading Attachments, page 30-16](#)
- [Testing for Viruses, page 30-16](#)

## Viewing Matched Content

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, except for DLP policy violation matches. You can also use the `$MatchedContent` action variable to include the matched content from message or content filter matches in the message subject.

If the attachment contains the matched content, the attachment's contents are displayed, as well as the reason it was quarantined, whether it was due to a DLP policy violation, content filter condition, message filter condition, or Image Analysis verdict.

When you view messages in the local quarantine that have triggered message or content filter rules, the GUI may display content that did not actually trigger the filter action (along with content that triggered the filter action). The GUI display should be used as a guideline for locating content matches, but does not necessarily reflect an exact list of content matches. This occurs because the GUI uses less strict content matching logic than is used in the filters. This issue applies only to the highlighting in the message body. The table that lists the matched strings in each part of the message, along with the associated filter rule, is correct.

**Figure 30-1** Matched Content Viewed in the Policy Quarantine

| Attachment Name | Matched Content                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Condition                           |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|
| FP1.1.txt       | <ul style="list-style-type: none"> <li>MS 38930 USA Facilities 662-646-0523 jsamuelson@acmecorp.com 7/17/06</li> <li>4929132070312710 Acme Corp Irene Gibbs 808 Sumner Street Greenwood</li> <li>MS 38930 USA Publishing 662-646-0522 igibbs@acmecorp.com 2/1/07</li> <li>4485231592071860 Acme Corp Kathy Lopez 808 Sumner Street</li> <li>Greenwood MS 38930 USA Marketing 662-646-0541 klopez@acmecorp.com</li> <li>2/1/07 4716298862510192 Acme Corp Marty Smith 808 Sumner Street</li> <li>Greenwood MS 38930 USA Engineering 662-646-0542</li> </ul> | DLP Classifier: Contact Information |

**Headers**

```
X-IronPort-AV: E=Sophos;i="4.43,282,1246818600";
d="txt?scan'208";a="178202"
Received: from d2.vmw023-bsd04.ibqa (HELO vmw023-bsd04.ibqa) ([172.22.107.1])
by c360q02.ibqa with ESMTP; 28 Jul 2009 16:25:03 +0530
Message-ID: <792087.518002035-sendEmail@vmw023-bsd04>
From: "user@test.com" <user@test.com>
To: "user1@test.com" <user1@test.com>
Subject: DLPTEST
Date: Tue, 28 Jul 2009 08:42:11 +0000
X-Mailer: sendEmail-1.55
MIME-Version: 1.0
Content-Type: multipart/mixed; boundary="-----MIME delimiter for sendEmail-538525.714612664"
```

**Message**

Test

| Name           | Size | Details                                |
|----------------|------|----------------------------------------|
| [message body] | 6    | ASCII text, with CRLF line terminators |
| FP1.1.txt      | 1K   | ASCII text                             |

## Downloading Attachments

You can download a message attachment by clicking the attachment's file name in the Message Parts or Matched Content section. AsyncOS displays a warning that attachments from unknown sources may contain viruses and asks you if you want to continue. Download attachments that may contain viruses at your own risk. You can also download the message body by clicking [message body] in the Message Parts section.

## Testing for Viruses

To test the message for viruses, click **Start Test**. Use a quarantine to hold messages until you are sure that your anti-virus signatures have been updated.

Testing for viruses sends a copy of the message to the anti-virus engine, not the message itself. The verdict from the anti-virus engine is returned and displayed above the Quarantines area.

## About Rescanning of Quarantined Messages

When a message is released from all queues in which it has been quarantined, the following rescanning occurs, depending on the features enabled for the appliance and for the mail policy that originally quarantined the message:

- Messages released from Policy and Virus quarantines are rescanned by the anti-virus engine.
- Messages released from the Outbreak quarantine are rescanned by the anti-spam and anti-virus engines. (For information about rescanning of messages while in the Outbreak quarantine, see [Chapter 14, “Outbreak Filters.”](#) )
- Messages released from the File Analysis quarantine are rescanned for threats.
- Messages with attachments are rescanned by the file reputation service upon release from Policy, Virus, and Outbreak quarantines.

Upon rescanning, if the verdict produced matches the verdict produced the previous time the message was processed, the message is not re-quarantined. Conversely, if the verdict is different, the message could be sent to another quarantine.

The rationale is to prevent messages from looping back to the quarantine indefinitely. For example, suppose a message is encrypted and therefore sent to the Virus quarantine. If an administrator releases the message, the anti-virus engine will still not be able to decrypt it; however, the message should not be re-quarantined or a loop will be created and the message will never be released from the quarantine. Since the two verdicts are the same, the system bypasses the Virus quarantine the second time.

## The Outbreak Quarantine

The Outbreak quarantine is present when a valid Outbreak Filters feature license key has been entered. The Outbreak Filters feature sends messages to the Outbreak quarantine, depending on the threshold set. For more information, see [Chapter 14, “Outbreak Filters.”](#)

The Outbreak quarantine functions just like other quarantines—you can search for messages, release or delete messages, and so on.

The Outbreak quarantine has some additional features not available in other quarantines: the Manage by Rule Summary link, the Send to Cisco feature when viewing message details, and the option to sort messages in search results by the Scheduled Exit time.

If the license for the Outbreak Filters feature expires, you will be unable to add more messages to the Outbreak quarantine. Once the messages currently in the quarantine have expired and the Outbreak quarantine becomes empty, it is no longer shown in the Quarantines listing in the GUI.

### Related Topics

- [Rescanning Messages in an Outbreak Quarantine, page 30-17](#)
- [Manage by Rule Summary Link, page 30-18](#)
- [Reporting False Positives or Suspicious Messages to Cisco Systems, page 30-18](#)

## Rescanning Messages in an Outbreak Quarantine

Messages placed in the Outbreak quarantine are automatically released if newly published rules deem the quarantined message no longer a threat.

If anti-spam and anti-virus are enabled on the appliance, the scanning engines scan every message released from the Outbreak quarantine based on the mail flow policy that applies to the message.

## Manage by Rule Summary Link

Click the Manage by Rule Summary link next to the Outbreak quarantine in the quarantine listing to view the Manage by Rule Summary page. You can perform message actions (Release, Delete, Delay Exit) on all of the messages in the quarantine based on which outbreak rule caused the message to be quarantined. This is ideal for clearing out large numbers of messages from the Outbreak quarantine. For more information, see the topics under [Outbreak Quarantine and the Manage by Rule Summary View](#), page 14-22.

## Reporting False Positives or Suspicious Messages to Cisco Systems

When viewing message details for a message in the Outbreak quarantine, you can send the message to Cisco to report false positives or suspicious messages.

### Procedure

---

- Step 1** Navigate to a message in the Outbreak quarantine.
  - Step 2** In the Message Details section, select the **Send a Copy to Cisco Systems** check box.
  - Step 3** Click **Send**.
-



## Spam Quarantine

---

- [Overview of the Spam Quarantine, page 31-1](#)
- [Local Versus External Spam Quarantine, page 31-1](#)
- [Setting Up the Local Spam Quarantine, page 31-2](#)
- [Using Safelists and Blocklists to Control Email Delivery Based on Sender, page 31-7](#)
- [Configuring Spam Management Features for End Users, page 31-14](#)
- [Managing Messages in the Spam Quarantine, page 31-22](#)
- [Disk Space for the Spam Quarantine, page 31-24](#)
- [About Disabling the Spam Quarantine, page 31-24](#)
- [Troubleshooting Spam Quarantine Features, page 31-24](#)

## Overview of the Spam Quarantine

The spam quarantine (also known as ISQ, End-User Quarantine, and EUQ) provides a safeguard mechanism for organizations that are concerned about “false positives” — that is, legitimate email messages that the appliance has deemed to be spam. When the appliance determines that a message is spam or suspected spam, you may want to let the recipient or an administrator review the message before delivering or deleting it. The spam quarantine stores messages for this purpose.

Administrative users of the Email Security appliance can view all messages in a spam quarantine. End users, usually the message recipients, can view their own quarantined messages in a slightly different web interface.

The spam quarantine is separate from policy, virus, and outbreak quarantines.

### Related Topics

- [Chapter 13, “Anti-Spam”](#)
- [Chapter 30, “Policy, Virus, and Outbreak Quarantines”](#)

## Local Versus External Spam Quarantine

A local spam quarantine stores spam and suspect spam on the Email Security appliance. An external spam quarantine can store these messages on a separate Cisco Content Security Management appliance.

Consider using an external spam quarantine if:

- You want a centralized location to store and manage spam from multiple Email Security appliances.
- You want to store more spam than the Email Security appliance can hold.
- You want to regularly back up the spam quarantine and its messages.

#### Related Topics

- [Disk Space for the Spam Quarantine, page 31-24](#)
- [Working with an External Spam Quarantine, page 42-2](#)

## Setting Up the Local Spam Quarantine

**Table 31-1**      *How to Send Messages to a Spam Quarantine*

|               | Do This                                                                                             | More Info                                                                                                                                                                                                                                                                                                        |
|---------------|-----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Enable the Anti-Spam feature if you have not yet done so.                                           | <a href="#">Chapter 13, “Anti-Spam”</a>                                                                                                                                                                                                                                                                          |
| <b>Step 2</b> | Enable and configure quarantine settings.                                                           | <a href="#">Enabling and Configuring the Spam Quarantine, page 31-3</a>                                                                                                                                                                                                                                          |
| <b>Step 3</b> | Adjust the disk space allocated to the spam quarantine                                              | <a href="#">Managing Disk Space, page 33-16</a>                                                                                                                                                                                                                                                                  |
| <b>Step 4</b> | Enable browser access to the quarantine.                                                            | <a href="#">Configuring the IP Interface for Browser Access to the Spam Quarantine, page 31-4</a>                                                                                                                                                                                                                |
| <b>Step 5</b> | Configure the Email Security appliance to send spam to the quarantine.                              | <ul style="list-style-type: none"> <li>• <a href="#">How to Configure the Appliance to Scan Messages for Spam, page 13-2</a></li> <li>• <a href="#">Configuring a Mail Policy to Quarantine Spam, page 31-5</a></li> <li>• <a href="#">Limiting Which Recipients Have Mail Quarantined, page 31-5</a></li> </ul> |
| <b>Step 6</b> | Specify a default character encoding for messages that do not have this information in the heading. | <a href="#">Ensuring That Message Text Displays Correctly, page 31-6</a>                                                                                                                                                                                                                                         |

#### Related Topics

- [Enabling and Configuring the Spam Quarantine, page 31-3](#)
- [Configuring the IP Interface for Browser Access to the Spam Quarantine, page 31-4](#)
- [Configuring Administrative User Access to the Spam Quarantine, page 31-4](#)
- [Configuring a Mail Policy to Quarantine Spam, page 31-5](#)
- [Limiting Which Recipients Have Mail Quarantined, page 31-5](#)
- [Ensuring That Message Text Displays Correctly, page 31-6](#)
- [Spam Quarantine Language, page 31-6](#)

## Enabling and Configuring the Spam Quarantine


**Note**

If you use an external spam quarantine, you will configure the settings described in this section on the Security Management appliance.

**Procedure**

- Step 1** Select **Monitor > Spam Quarantine**.
- Step 2** If you have not previously enabled the spam quarantine, select **Enable Spam Quarantine**.  
If you are editing spam quarantine settings, click the **Spam Quarantine** link in the Quarantine Name column of the Spam Quarantine section.
- Step 3** Specify options:

| Option                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Quarantine Size                   | <p>If you deselect <b>When storage space is full, automatically delete oldest messages first</b>, newer messages will not be added to a full quarantine. Cisco recommends that you enable this option so that a full quarantine will not cause messages to queue (back up) on your appliance.</p> <p>To manage disk space for your quarantine, see <a href="#">Managing Disk Space, page 33-16</a>.</p>                                                                                                                                                                                                                                                                                                               |
| Schedule Delete After             | <p>Specify the number of days to hold messages before deleting them.</p> <p>Cisco recommends that you configure the quarantine to delete older messages to prevent the quarantine from filling to capacity, but you can elect not to schedule automatic deletion.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Notify Cisco Upon Message Release | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Spam Quarantine Appearance        | <p><b>Logo</b></p> <p>By default, the Cisco logo is displayed at the top of the spam quarantine page when the user logs in to view quarantined messages.</p> <p>To use a custom logo instead, upload the logo. The logo should be a .jpg, .gif, or .png file that is at most 50 pixels high by 500 pixels wide.</p> <p><b>Login page message</b></p> <p>(Optional) Specify a login page message. This message is shown to end users and administrators when they log in to view the quarantine.</p> <p>If you do not specify a message, the following message appears:</p> <p style="padding-left: 40px;">Enter your login information below. If you are unsure what to enter, please contact your administrator.</p> |
| Administrative Users              | See <a href="#">Configuring Administrative User Access to the Spam Quarantine, page 31-4</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

- Step 4** Submit and commit your changes.



**What To Do Next**

- Return to [Setting Up the Local Spam Quarantine, page 31-2](#).

## Configuring the IP Interface for Browser Access to the Spam Quarantine

When administrators and end users access the spam quarantine, a separate browser window opens.

**Procedure**

- 
- Step 1** Choose **Network > IP Interfaces**.
- Step 2** Click the interface name (for this example, we will use the Management interface).
- Step 3** In the Spam Quarantine section, configure settings for access to the spam quarantine:
- By default, HTTP uses port 82 and HTTPS uses port 83.
  - Specify the URL that appears in notifications and in the spam quarantine browser window.  
If you do not want to expose the hostname of your Security Management appliance to end users, you can specify an alternate hostname.
- Step 4** Submit and commit your changes.
- 

**What To Do Next**

Ensure that your DNS server can resolve the hostname that you specified for spam quarantine access.

## Configuring Administrative User Access to the Spam Quarantine

All users with administrator privileges can change spam quarantine settings and view and manage messages in the spam quarantine. You do not need to configure spam quarantine access for administrator users.

If you configure access to the spam quarantine for users with the following roles, they can view, release, and delete messages in the spam quarantine:

- Operator
- Read-only operator
- Help desk user
- Guest
- Custom user roles that have spam quarantine privileges

These users cannot access spam quarantine settings.

**Before You Begin**

Create users or custom user roles that have access to the spam quarantine. For more information, see [Chapter 32, “Distributing Administrative Tasks.”](#)

### Procedure

- 
- Step 1** If you are not already editing the spam quarantine settings page:
- Select **Monitor > Spam Quarantine**.
  - Click the **Spam Quarantine** link in the Quarantine Name column of the Spam Quarantine section.
- Step 2** Click the link for the type of user to add: local, externally authenticated, or custom role.
- If you have already added users or roles, click a username or role to view all eligible users or roles.
- Step 3** Select the users or roles that you want to add.
- Users with Administrator privileges are not listed because they automatically have full access to the spam quarantine.
- Step 4** Click **OK**.
- Step 5** Submit and commit your changes.
- 

### Related Topics

- [Configuring End-User Access to the Spam Quarantine, page 31-17](#)

## Configuring a Mail Policy to Quarantine Spam

Once you have enabled the spam quarantine, you can configure a mail policy to send spam or suspected spam to that quarantine. Anti-spam scanning must be enabled on the mail policy in order to send mail to the spam quarantine.

For complete information about configuring mail policies to quarantine spam, see [Defining Anti-Spam Policies, page 13-7](#).

### Procedure

- 
- Step 1** On the Mail Policies > Incoming Mail Policies page, click the link in the Anti-Spam column for the corresponding mail policy.
- Step 2** In the Anti-Spam Settings section, select **Use IronPort Anti-Spam service**.
- Step 3** In the Positively-Identified Spam Settings section, select **Spam Quarantine** for the **Apply This Action to Message** option.
- Step 4** Configure settings for suspected spam and marketing email.
- Step 5** Submit and commit your changes.
- 

## Limiting Which Recipients Have Mail Quarantined

You can use multiple mail policies (Mail Policies > Incoming Mail Policy) to specify a list of recipient addresses for which mail will not be quarantined. Select 'Deliver' or 'Drop' instead of quarantine when configuring the anti-spam settings for the mail policy.

## Ensuring That Message Text Displays Correctly

AsyncOS attempts to determine the character set of a message based on the encoding that is specified in the message headers. However, if the encoding specified in the headers does not match that of the actual text, the message will not be displayed properly when viewed in the spam quarantine. This situation is more likely to occur with spam messages.

To ensure that message text displays correctly for these messages, see [Specifying a Default Encoding, page 31-6](#).

### Related Topics

- [Specifying a Default Encoding, page 31-6](#)

## Specifying a Default Encoding

If an incoming message does not have a character set encoding specified in the headers, you can configure your appliance to specify a default encoding.

Doing so will help ensure that these types of messages display properly in the spam quarantine. However, specifying a default encoding can cause messages in other character sets to display incorrectly. This setting applies only to messages that do not specify the encoding in the message headers. Generally, you would only set a default encoding if you expect the majority of your mail that falls into this category to be of one specific encoding.

For example, if most quarantined messages that do not specify the character set encoding in the message headers are in Japanese (ISO-2022-JP), you can set the encoding on the Scan Behavior page as **Japanese (ISO-2022-JP)**.

### Procedure

- 
- |               |                                                                                                          |
|---------------|----------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Click <b>Security Services &gt; Scan Behavior</b> .                                                      |
| <b>Step 2</b> | Under Global Settings, click <b>Edit Global Settings</b> .                                               |
| <b>Step 3</b> | From the <b>Encoding to use when none is specified</b> drop-down list, select the desired encoding type. |
| <b>Step 4</b> | Click <b>Submit</b> .                                                                                    |
| <b>Step 5</b> | Click <b>Commit Changes</b> .                                                                            |
- 

## Spam Quarantine Language

Each user selects a language in the spam quarantine from the Options menu at the top right of the window.

# Using Safelists and Blocklists to Control Email Delivery Based on Sender

Administrators and end users can use safelists and blocklists to help determine which messages are spam. Safelists specify senders and domains that are never treated as spam. Blocklists specify senders and domains that are always treated as spam.

You can allow end users (email users) to manage the safelist and blocklist for their own email accounts. 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 being sent to his inbox. On the other hand, end users may find that emails from specific senders are sent to their spam quarantine when they do not want them to be treated as spam. To ensure that messages from these senders are not quarantined, they may want to add the senders to their safelists.

Changes that end users and administrators make are visible to and can be changed by either.

## Related Topics

- [Message Processing of Safelists and Blocklists, page 31-7](#)
- [Enabling Safelists and Blocklists, page 31-8](#)
- [External Spam Quarantine and Safelist/Blocklists, page 31-8](#)
- [Adding Senders and Domains to Safelists and Blocklists \(Administrators\), page 31-9](#)
- [About End-User Access to Safelists and Blocklists, page 31-11](#)
- [Synchronizing Safelists/Blocklists on Multiple Email Security Appliances \(Deployments Without a Security Management Appliance\), page 31-12](#)
- [Backing Up and Restoring the Safelist/Blocklist, page 31-13](#)
- [Troubleshooting Safelists and Blocklists, page 31-13](#)

## Message Processing of Safelists and Blocklists

A sender's being on a safelist or blocklist does not prevent the appliance from scanning a message for viruses or determining if the message meets the criteria for a content-related mail policy. Even if the sender of a message is on the recipient's safelist, the message may not be delivered to the end user depending on other scanning settings and results.

When you enable safelists and blocklists, the appliance scans the messages against the safelist/blocklist database immediately before anti-spam scanning. If the appliance detects a sender or domain that matches a safelist or blocklist entry, the message will be splintered if there are multiple recipients (and the recipients have different safelist/blocklist settings). For example, a message is sent to both recipient A and recipient B. Recipient A has safelisted the sender, whereas recipient B does not have an entry for the sender in the safelist or the blocklist. In this case, the message may be split into two messages with two message IDs. The message sent to recipient A is marked as safelisted with an *X-SLBL-Result-Safelist* header and skips anti-spam scanning, whereas the message bound for recipient B is scanned by the anti-spam scanning engine. Both messages then continue along the pipeline (through anti-virus scanning, content policies, and so on) and are subject to any configured settings.

If a message sender or domain is blocklisted, the delivery behavior depends on the blocklist action that you specify when you enable the safelist/blocklist feature. Similar to safelist delivery, the message is splintered if there are different recipients with different safelist/blocklist settings. The blocklisted message splinter is then quarantined or dropped, depending on the blocklist action settings. If the

blocklist action is configured to quarantine, the message is scanned and eventually quarantined. If the blocklist action is configured to delete, the message is dropped immediately after safelist/blocklist scanning.

Because safelists and blocklists are maintained in the spam quarantine, delivery behavior is also contingent on other anti-spam settings. For example, if you configure the “Accept” mail flow policy in the Host Access Table (HAT) to skip anti-spam scanning, then users who receive mail on that listener will not have their safelist and blocklist settings applied to mail received on that listener. Similarly, if you create a mail flow policy that skips anti-spam scanning for certain message recipients, these recipients will not have their safelist and blocklist settings applied.

#### Related Topics

- [Enabling Safelists and Blocklists, page 31-8](#)
- [External Spam Quarantine and Safelist/Blocklists, page 31-8](#)

## Enabling Safelists and Blocklists

#### Before You Begin

- The spam quarantine must be enabled. See [Setting Up the Local Spam Quarantine, page 31-2](#).

#### Procedure

---

**Step 1** Select **Monitor > Spam Quarantine**.

**Step 2** In the **End-User Safelist/Blocklist (Spam Quarantine)** section, select **Enable**.

**Step 3** Select **Enable End User Safelist/Blocklist Feature**.

**Step 4** Select **Quarantine** or **Delete** for the Blocklist Action.

**Step 5** Specify the **Maximum List Items Per User**.

This is the maximum number of addresses or domains for each list, for each recipient. If you allow a large number of list entries per user, system performance might be adversely affected.

**Step 6** Submit and commit your changes.

---

## External Spam Quarantine and Safelist/Blocklists

If you use an external spam quarantine on a Security Management appliance, the safelist/blocklist is saved on the management appliance. This gives you a single location to manage safe and blocked senders for all appliances.

Because the Email Security appliance evaluates senders in safelists and blocklists when processing incoming mail, safelists and blocklists that are stored on a Security Management appliance must be sent to the Email Security appliance in order to be applied to incoming mail. When you configure the safelist/blocklist feature on a Security Management appliance, you configure the frequency of these updates.

For more information about working with external safelists and blocklists on a Security Management appliance, see the topics under [Working with an External Spam Quarantine, page 42-2](#), and the *Cisco Content Security Management Appliance User Guide*.

## Adding Senders and Domains to Safelists and Blocklists (Administrators)

Manage safelists and blocklists via the spam quarantine interface.

You can also see whether many recipients (end users in your organization) have whitelisted or blacklisted a particular sender or domain.

Administrators see and work with the superset of the same entries that each end user sees and works with.

### Before You Begin

- Ensure that you can access the spam quarantine. See [Accessing the Spam Quarantine \(Administrative Users\)](#), page 31-22.
- Enable access to the safelist/blocklist. See [Enabling Safelists and Blocklists](#), page 31-8.
- (Optional) To import safelist/blocklists instead of building these lists using the procedure in this section, use the process described in [Backing Up and Restoring the Safelist/Blocklist](#), page 31-13.
- Understand the required format of safelist and blocklist entries. See [Syntax for Safelists and Blocklist Entries](#), page 31-10.

### Procedure

---

- Step 1** Using your browser, access the spam quarantine.
- Step 2** Log in.
- Step 3** Select the **Options** drop-down menu in the upper right corner of the page.
- Step 4** Choose **Safelist** or **Blocklist**.
- Step 5** (Optional) Search for a sender or recipient.
- Step 6** Do one or more of the following:

| To                                                                                                   | Do This                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Add multiple senders for a recipient                                                                 | <ol style="list-style-type: none"> <li>1. Select <b>View by: Recipient</b></li> <li>2. Click <b>Add</b>, or click <b>Edit</b> for a recipient.</li> <li>3. Enter or edit the recipient email address.</li> <li>4. Enter sender email addresses and domains.<br/>Put each entry on a separate line, or separate each entry with a comma.</li> <li>5. Click <b>Submit</b>.</li> </ol> |
| Add multiple recipients for a sender                                                                 | <ol style="list-style-type: none"> <li>1. Select <b>View by: Sender</b></li> <li>2. Click <b>Add</b>, or click <b>Edit</b> for a sender.</li> <li>3. Enter or edit the sender address or domain.</li> <li>4. Enter recipient email addresses.<br/>Put each entry on a separate line, or separate each entry with a comma.</li> <li>5. Click <b>Submit</b>.</li> </ol>               |
| Delete all senders associated with a recipient<br><br>Delete all recipients associated with a sender | <ol style="list-style-type: none"> <li>1. Select a <b>View by</b> option.</li> <li>2. Click a trash can icon to delete an entire table row.</li> </ol>                                                                                                                                                                                                                              |
| Delete individual senders for a recipient<br><br>Delete individual recipients for a sender           | <ol style="list-style-type: none"> <li>1. Select a <b>View by</b> option.</li> <li>2. Click <b>Edit</b> for an individual recipient or sender.</li> <li>3. Add or remove entries from the text box. You must leave at least one entry.</li> <li>4. Click <b>Submit</b>.</li> </ol>                                                                                                  |

**Related Topics**

- [Syntax for Safelists and Blocklist Entries, page 31-10](#)
- [Clearing All Safelists and Blocklists, page 31-11](#)

**Syntax for Safelists and Blocklist Entries**

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

- user@domain.com
- server.domain.com
- domain.com
- [10.1.1.0]
- [ipv6:2001:DB8:1::1]
- user@[1.2.3.4]

- `user@[ipv6:2001:db8::1]`

An identical entry, such as a sender address or a domain, cannot be included on both the safelist and the blocklist at the same time. However, a domain can be on a safelist while an email address for a sender belonging to that domain is on the blocklist (or vice versa), and both rules apply. For example, if *example.com* is on the safelist, *george@example.com* can be on the blocklist. In this case, the appliance delivers all mail from *example.com* without scanning for spam, except mail from *george@example.com*, which is treated as spam.

It is not possible allow or block a range of subdomains using the following syntax: *.domain.com*. However, it is possible to block a specific domain using the following syntax: *server.domain.com*.

## Clearing All Safelists and Blocklists

If you need to delete all safelist and blocklist entries, including all senders and all recipients, import a file with no entries using the procedure in [Backing Up and Restoring the Safelist/Blocklist](#), page 31-13.

## About End-User Access to Safelists and Blocklists

End users access their safelist and blocklist via the spam quarantine. To configure end-user access to the spam quarantine, see [Setting Up End-User Access to the Spam Quarantine via Web Browser](#), page 31-16.

You may want to give your end users the URL of the spam quarantine and the instructions below, as applicable.

### Related Topics

- [Adding Entries to Safelists \(End Users\)](#), page 31-11
- [Adding Senders to Blocklists \(End Users\)](#), page 31-12

## Adding Entries to Safelists (End Users)



### Note

Delivery of messages from safelisted senders depends on other settings that are configured in the system. See [Message Processing of Safelists and Blocklists](#), page 31-7.

End users can add senders to safelists in two ways:

- [Adding the Sender of a Quarantined Message to the Safelist](#), page 31-11
- [Adding Senders to the Safelist Without a Quarantined Message](#), page 31-12

### Adding the Sender of a Quarantined Message to the Safelist

End users can add senders to the safelist if the message has been sent to the spam quarantine.

#### Procedure

- Step 1** From the spam quarantine, select the checkbox next to the message.
- Step 2** Choose **Release and Add to Safelist** from the drop-down menu.



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 Senders to the Safelist Without a Quarantined Message

#### Procedure

---

- Step 1** Access the spam quarantine via browser.
  - Step 2** Select the **Options** drop-down menu in the upper right corner of the page.
  - Step 3** Choose **Safelist**.
  - Step 4** From the Safelist dialog box, enter the email address or domain. You can enter multiple domains and email addresses, separated by commas.
  - Step 5** Click **Add to List**.
- 

### Adding Senders to Blocklists (End Users)

Messages from blocklisted senders may be rejected or quarantined, depending on the safelist/blocklist action settings defined by your administrator.



#### Note

You can add blocklist entries only using this procedure.

---

#### Procedure

---

- Step 1** Log in to the spam quarantine.
  - Step 2** Select the **Options** drop-down menu in the upper right corner of the page.
  - Step 3** Enter the domain or email address that you want to blocklist. You can enter multiple domains and email addresses, separated by commas.
  - Step 4** Click **Add to List**.
- 

## Synchronizing Safelists/Blocklists on Multiple Email Security Appliances (Deployments Without a Security Management Appliance)

If you use multiple Email Security appliances without a Security Management appliance, you may need to manually synchronize the safelist/blocklist and its configuration settings across the different Email Security appliances.

You can export and import a .csv file using the procedure described in [Backing Up and Restoring the Safelist/Blocklist, page 31-13](#), then use FTP to upload and download the file. For more information, see [Appendix A, "FTP, SSH, SCP, and Telnet Access."](#)

## Backing Up and Restoring the Safelist/Blocklist

Before you upgrade your appliance or run the installation wizard, you should back up the safelist/blocklist database. Safelist/blocklist information is not included in the main XML configuration file that contains your appliance configuration settings.

You can also use this procedure to save a copy of the safelist/blocklist to synchronize multiple Email Security appliances.

### Procedure

- Step 1** Select **System Administration > Configuration File**.
- Step 2** Scroll to the **End-User Safelist/Blocklist Database (Spam Quarantine)** section.

| To                            | Do This                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Export the safelist/blocklist | <p>Note the path and filename of the .csv file, and modify as needed.</p> <p>Click <b>Backup Now</b>.</p> <p>The appliance saves a .csv file to the <code>/configuration</code> directory of the appliance using the following naming convention:</p> <p><i>slbl&lt;serial number&gt;&lt;timestamp&gt;.csv</i></p>                                                                                                                              |
| Import the safelist/blocklist | <p> <b>Caution</b> This process will overwrite all existing entries in safelists and blocklists for all users.</p> <p>Click <b>Select File to Restore</b>.</p> <p>Select the desired file from the list of files in your configuration directory.</p> <p>Select the safelist/blocklist backup file that you want to restore.</p> <p>Click <b>Restore</b>.</p> |

## Troubleshooting Safelists and Blocklists

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 ISQ\_log files or the antispam log files. Emails that are safelisted are marked as safelisted with an *X-SLBL-Result-Safelist* header. Emails that are blocklisted are marked as blocklisted with an *X-SLBL-Result-Blocklist* header.

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

For more information about alerts, see [Chapter 33, “Alerts.”](#)

For more information about log files, see [Chapter 38, “Logging.”](#)

**Related Topics**

- [Message from Safelisted Sender Was Not Delivered, page 31-14](#)

## Message from Safelisted Sender Was Not Delivered

**Problem** Message from a safelisted sender was not delivered.

**Solution** Possible causes:

- The message was dropped for malware or content violations. See [Message Processing of Safelists and Blocklists, page 31-7](#).
- If you have multiple appliances and the sender was recently added to the safelist, then safelist/blocklists might not have been synchronized at the time the message was processed. See [External Spam Quarantine and Safelist/Blocklists, page 31-8](#) and [Synchronizing Safelists/Blocklists on Multiple Email Security Appliances \(Deployments Without a Security Management Appliance\), page 31-12](#).

## Configuring Spam Management Features for End Users

| To                                                                                                                                                                     | See                                                                                                 |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| Understand the benefits and limitations of the different authentication methods for end-user access to spam management features.                                       | <a href="#">Configuring End-User Access to the Spam Quarantine, page 31-17</a> and subsections      |
| Allow end users to access the spam quarantine directly via browser.                                                                                                    | <a href="#">Authentication Options for End Users Accessing Spam Management Features, page 31-15</a> |
| Send users a notification when messages addressed to them are routed to the spam quarantine.<br><br>Notifications can include links for access to the spam quarantine. | <a href="#">Notifying End Users About Quarantined Messages, page 31-19</a>                          |
| Allow users to specify email addresses and domains of senders whom they know to be safe, and of senders whom they know to be sending spam or other unwanted mail.      | <a href="#">Using Safelists and Blocklists to Control Email Delivery Based on Sender, page 31-7</a> |

**Related Topics**

- [Authentication Options for End Users Accessing Spam Management Features, page 31-15](#)
- [Setting Up End-User Access to the Spam Quarantine via Web Browser, page 31-16](#)
- [Notifying End Users About Quarantined Messages, page 31-19](#)

## Authentication Options for End Users Accessing Spam Management Features


**Note**

Mailbox authentication does not allow users to view messages addressed to an email alias.

| For End-User Spam Quarantine Access                                                                                   | Do This                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Directly via web browser, authentication required<br>and<br>Via a link in a notification, authentication required     | <ol style="list-style-type: none"> <li>1. In the End User Quarantine Access settings, choose <b>LDAP</b> or <b>Mailbox (IMAP/POP)</b>.</li> <li>2. In the Spam Notifications settings, deselect <b>Enable login without credentials for quarantine access</b>.</li> </ol> |
| Directly via web browser, authentication required<br>and<br>Via a link in a notification, authentication not required | <ol style="list-style-type: none"> <li>1. In the End User Quarantine Access settings, choose <b>LDAP</b> or <b>Mailbox (IMAP/POP)</b>.</li> <li>2. In the Spam Notifications settings, select <b>Enable login without credentials for quarantine access</b>.</li> </ol>   |
| Only via a link in a notification, authentication not required                                                        | In the End User Quarantine Access settings, choose <b>None</b> as the authentication method.                                                                                                                                                                              |
| No access                                                                                                             | In the End User Quarantine Access settings, deselect <b>Enable End-User Quarantine Access</b> .                                                                                                                                                                           |

**Related Topics**

- [LDAP Authentication Process, page 31-15](#)
- [IMAP/POP Authentication Process, page 31-16](#)
- [Configuring End-User Access to the Spam Quarantine, page 31-17](#)
- [Notifying End Users About Quarantined Messages, page 31-19](#)
- [Authenticating End-Users of the Spam Quarantine, page 25-43](#)
- [About End-User Access to Safelists and Blocklists, page 31-11](#)

### LDAP Authentication Process

1. A user enters his or her username and password into the web UI login page.
2. The spam quarantine connects to the specified LDAP server either to perform an anonymous search or as an authenticated user with the specified “Server Login” DN and password. For Active Directory, you will usually need to have the server connect on the “Global Catalog port” (it is in the 6000s) and you need to create a low privilege LDAP user that the spam quarantine can bind as in order to execute the search.
3. The spam quarantine then searches for the user using the specified BaseDN and Query String. When a user’s LDAP record is found, the spam quarantine then extracts the DN for that record and attempts bind to the directory using the user records’ DN and the password they entered originally. If this password check succeeds then the user is properly authenticated, but the spam quarantine still needs to determine which mailboxes’ contents to show for that user.

4. Messages are stored in the spam quarantine using the recipient's envelope address. After a user's password is validated against LDAP, the spam quarantine then retrieves the “Primary Email Attribute” from the LDAP record to determine which envelope address they should show quarantined messages for. The “Primary Email Attribute” can contain multiple email addresses which are then used to determine what envelope addresses should be displayed from the quarantine for the authenticated user.

#### Related Topics

- [Authenticating End-Users of the Spam Quarantine, page 25-43](#)

## IMAP/POP Authentication Process

1. Depending on your mail server configuration, a user enters their username (joe) or email address (joe@example.com) and password into the web UI login page. You can modify the Login Page Message to tell your users whether they should enter a full email address or just their username (see [Configuring End-User Access to the Spam Quarantine, page 31-17](#)).
2. The spam quarantine connects to the IMAP or POP server and uses the entered login (either username or email address) and password to try to log into the IMAP/POP server. If the password is accepted then the user is considered authenticated and the spam quarantine immediately logs out of the IMAP/POP server.
3. Once the user is authenticated, the spam quarantine lists email for the user, based on the email address:
  - If you have configured the spam quarantine to specify a domain to append to bare usernames (like joe), then this domain is appended and that fully qualified email address is used to search for matching envelopes in the quarantine.
  - Otherwise, the spam quarantine uses the entered email address to search for matching envelopes.

For more information about IMAP, see the University of Washington web site:

<http://www.washington.edu/imap/>

## Setting Up End-User Access to the Spam Quarantine via Web Browser

|               | Do This                                                                                                                                                                                                                                   | More Information                                                                                    |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Understand the benefits and limitations of the different authentication methods for end-user access to spam management features.                                                                                                          | <a href="#">Authentication Options for End Users Accessing Spam Management Features, page 31-15</a> |
| <b>Step 2</b> | If you will authenticate end users using LDAP, configure an LDAP server profile, including the <b>Spam Quarantine End-User Authentication Query</b> settings on the <b>System Administration &gt; LDAP &gt; LDAP Server Profile</b> page. | <a href="#">Authenticating End-Users of the Spam Quarantine, page 25-43</a> and subsections         |

|               | Do This                                                       | More Information                                                                           |
|---------------|---------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| <b>Step 3</b> | Configure end-user access to the spam quarantine.             | <a href="#">Configuring End-User Access to the Spam Quarantine, page 31-17</a>             |
| <b>Step 4</b> | Determine the URL for end-user access to the spam quarantine. | <a href="#">Determining the URL for End-User Access to the Spam Quarantine, page 31-18</a> |

**Related Topics**

- [Configuring End-User Access to the Spam Quarantine, page 31-17](#)
- [Determining the URL for End-User Access to the Spam Quarantine, page 31-18](#)
- [Which Messages an End User Sees, page 31-18](#)

## Configuring End-User Access to the Spam Quarantine

Administrative users can access the spam quarantine whether or not end-user access is enabled.

**Before You Begin**

See requirements in [Authentication Options for End Users Accessing Spam Management Features, page 31-15](#).

**Procedure**

- 
- Step 1** Select **Monitor > Spam Quarantine**.
- Step 2** Click the **Spam Quarantine** link in the Quarantine Name column of the Spam Quarantine section.
- Step 3** Scroll down to the **End-User Quarantine Access** section.
- Step 4** Select **Enable End-User Quarantine Access**.
- Step 5** Specify the method to use to authenticate end users when they attempt to view their quarantined messages.

| Select This Option | More Information |
|--------------------|------------------|
| None               | —                |

| Select This Option | More Information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Mailbox (IMAP/POP) | <p>For sites without an LDAP directory to use for authentication, the quarantine can validate user email addresses and passwords against a standards-based IMAP or POP server that holds their mailbox.</p> <p>When logging in to the spam quarantine, end users enter their full email address and mailbox password.</p> <p>If the POP server advertises APOP support in the banner, then for security reasons (i.e., to avoid sending the password in the clear) the Cisco appliance will only use APOP. If APOP is not supported for some or all users then the POP server should be reconfigured to not advertise APOP.</p> <p>Select SSL if you have configured your server to use it. If users enter username only, you can specify a domain to add to automatically complete the email address. Enter the domain of the envelope for users logging in to “Append Domain to Unqualified Usernames.”</p> |
| LDAP               | Configure LDAP settings as described in the sections referenced in the Before You Begin section of this topic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

**Step 6** Specify whether or not to display message bodies before messages are released.

If this box is selected, users may not view the message body via the spam quarantine page. Instead, to view the body of a quarantined message, users must release the message and view it in their mail application (such as Microsoft Outlook). You can use this feature for policy and regulation compliance — for example, if a regulation requires that all viewed email be archived.

**Step 7** Submit and commit your changes.

#### What To Do Next

(Optional) Customize the page that users see when they access the spam quarantine, if you have not yet done so. See setting descriptions in [Enabling and Configuring the Spam Quarantine, page 31-3](#).

## Determining the URL for End-User Access to the Spam Quarantine

The URL that end users can use to directly access the spam quarantine is formed from the hostname of the machine and the settings (HTTP/S and port numbers) configured on the IP interface on which the quarantine has been enabled. For example, `HTTP://mail3.example.com:82`.

## Which Messages an End User Sees

Generally, end users see only their own messages in the spam quarantine.

Depending on the method of access (via notification or directly via web browser) and authentication method (LDAP or IMAP/POP), users may see mail for multiple email addresses in the spam quarantine.

When LDAP authentication is used, if the Primary Email attribute has multiple values in the LDAP directory, all of those values (addresses) will be associated with the user. Therefore, quarantined messages addressed to all email addresses associated with the end user in the LDAP directory are present in the quarantine.

If the authentication method is IMAP/POP, or the user accesses the quarantine directly via a notification, then the quarantine will display only messages for that user's email address (or the address to which the notification was sent).

For information about messages that are sent to aliases of which the user is a member, see [Recipient Email Mailing List Aliases and Spam Notifications, page 31-20](#).

#### Related Topics

- [Configuring End-User Access to the Spam Quarantine, page 31-17](#)
- [Recipient Email Mailing List Aliases and Spam Notifications, page 31-20](#)

## Notifying End Users About Quarantined Messages

You can configure the system to send a notification email to some or all users when they have spam and suspected spam messages in the spam quarantine.

By default, spam notifications list the user's quarantined messages. Notifications can also include a link that users can click in order to view their quarantined messages in the spam quarantine. These links do not expire. The user can view the quarantined messages and decide whether to have them delivered to their inbox or delete them.



#### Note

In cluster configurations, you can choose which users receive notifications only at the machine level.

#### Before You Begin

- For end users to manage messages listed in notifications, they must be able to access the spam quarantine. See [Configuring End-User Access to the Spam Quarantine, page 31-17](#).
- Understand the authentication options for managing spam using notifications. See [Authentication Options for End Users Accessing Spam Management Features, page 31-15](#).
- If end users receive email at multiple aliases, see [Recipient Email Mailing List Aliases and Spam Notifications, page 31-20](#).

#### Procedure

- Step 1** Select **Monitor > Spam Quarantine**.
- Step 2** Click the **Spam Quarantine** link in the Quarantine Name column of the Spam Quarantine section.
- Step 3** Scroll down to the **Spam Notifications** section.
- Step 4** Select **Enable Spam Notification**.
- Step 5** Specify options.
 

To customize the message body:

  - a. (Optional) Customize the default text and variables.
 

The following message variables are expanded to the actual value for the specific end user:

    - **New Message Count** (%new\_message\_count%) — The number of new messages since the user last logged in.
    - **Total Message Count** (%total\_message\_count%) — The number of messages for the user in the spam quarantine.



- **Days Until Message Expires** (%days\_until\_expire%)
- **Quarantine URL** (%quarantine\_url%) — URL to log in to the quarantine and view messages.
- **Username** (%username%)
- **New Message Table** (%new\_quarantine\_messages%) — A list of the user's new messages in the quarantine.

To insert a variable, place the cursor where you would like the variable inserted and then click the name of the variable in the Message Variables listing on the right. Or type in the variable.

- If you have enabled an authentication method in the End User Quarantine Access section on this page:
  - To automatically log users in to the spam quarantine when they access it by clicking a link in a notification, select **Enable login without credentials for quarantine access**. End users can release messages simply by clicking “Release” links in the notification.
  - To require users to log in to the spam quarantine when they access it by clicking a link in a notification, deselect this option. End users cannot release messages simply by clicking “Release” links in the notification.
- Click **Preview Message** to verify that the message is as you want it to be.

**Step 6** Submit and commit your changes.

### What To Do Next

To ensure that end users receive these notifications, consider recommending that they add the From: address for the spam quarantine notification emails to the “whitelist” in the junk mail settings of their mail application (such as Microsoft Outlook or Mozilla Thunderbird.)

### Related Topics

- [Recipient Email Mailing List Aliases and Spam Notifications, page 31-20](#)
- [Testing Notifications, page 31-21](#)
- [Troubleshooting Spam Notifications, page 31-21](#)

## Recipient Email Mailing List Aliases and Spam Notifications

Notifications can be sent to each Envelope Recipient that has quarantined email, including mailing lists and other aliases. Each mailing list receives a single digest. If you send notifications to a mailing list, all subscribers to the list will receive the notification. Users who belong to multiple email aliases, or who belong to LDAP groups that receive notifications, or who use several email addresses, may receive multiple spam notifications. The following table shows example situations in which users may receive multiple notifications.

**Table 31-2** Notifications per Address/Alias

| User | Email Addresses | Aliases | Notifications |
|------|-----------------|---------|---------------|
| Sam  | sam@example.com | —       | 1             |

**Table 31-2**      **Notifications per Address/Alias**

| User | Email Addresses                       | Aliases                                             | Notifications |
|------|---------------------------------------|-----------------------------------------------------|---------------|
| Mary | mary@example.com                      | dev@example.com<br>qa@example.com<br>pm@example.com | 4             |
| Joe  | joe@example.com,<br>admin@example.com | hr@example.com                                      | 3             |

If you use LDAP authentication, you can choose not to send notifications to mailing list aliases. Or, if you choose to send spam notifications to mailing list aliases, you can prevent some occurrences of multiple notifications. See [Spam Quarantine Alias Consolidation Queries, page 25-44](#).

Users who access the spam quarantine by clicking a link in a notification will not see quarantined messages for any other aliases that the end-user may have, unless the appliance is using a spam quarantine alias consolidation query for email notifications. If the notification was sent to a distribution list that is expanded after processing by the appliance, then multiple recipients may have access to the same quarantine for that list.

This means that all subscribers to a mailing list will receive the notification and can log in to the quarantine to release or delete messages. In this case, end users visiting the quarantine to view messages mentioned in a notification may find that those messages have already been deleted by other users.

**Note**

If you do not use LDAP and you do not want your end users to receive multiple email notifications, consider disabling notifications and instead allow end users to access the quarantine directly and authenticate via LDAP or POP/IMAP.

## Testing Notifications

You can test notifications by configuring a testing mail policy, and having spam quarantined for just a single user. Then, configure the spam quarantine notification settings: Select the **Enable Spam Notification** checkbox and do not select **Enable End-User Quarantine Access**. Then only the administrator configured in the **Deliver Bounced Messages To** field is notified of new spam in the quarantine.

## Troubleshooting Spam Notifications

### Related Topics

- [User Receives Multiple Notifications, page 31-21](#)
- [Recipient Does Not Receive Notifications, page 31-22](#)

### User Receives Multiple Notifications

**Problem** A user receives multiple spam notifications for a single message.

**Solution** Possible causes:

- The user has multiple email addresses and the spam message was sent to more than one of those addresses.

- The user is a member of one or more email aliases that received the spam message. To minimize duplications, and for more information, see [Recipient Email Mailing List Aliases and Spam Notifications, page 31-20](#).

### Recipient Does Not Receive Notifications

**Problem** Recipient is not receiving spam notifications.

**Solution**

- If notifications are being sent to the “Deliver Bounce Messages To:” address instead of to spam recipients, this means that spam notifications are enabled, but spam quarantine access is not enabled. See [Authentication Options for End Users Accessing Spam Management Features, page 31-15](#).
- Have the user check the junk mail settings of their email client.

## Managing Messages in the Spam Quarantine

This section explains how to work with messages in local or external spam quarantines.

Administrative users can see and manage all messages in the spam quarantine.

**Related Topics**

- [Accessing the Spam Quarantine \(Administrative Users\), page 31-22](#)
- [Searching for Messages in the Spam Quarantine, page 31-22](#)
- [Viewing Messages in the Spam Quarantine, page 31-23](#)
- [Delivering Messages in the Spam Quarantine, page 31-23](#)
- [Deleting Messages from the Spam Quarantine, page 31-24](#)

### Accessing the Spam Quarantine (Administrative Users)

- 
- Step 1** Select **Monitor > Spam Quarantine**, then click the number in the Messages column.  
The spam quarantine opens in a separate browser window.
- 

### Searching for Messages in the Spam Quarantine

**Procedure**

- 
- Step 1** Specify an envelope recipient.



**Note** You can enter a partial address.

---

- Step 2** Select whether the search results should match the exact recipient you entered, or whether the results should contain, start with, or end with your entry.

- Step 3** Enter a date range to search through. Click the calendar icons to select a date.
- Step 4** Specify a From: address, and select whether the search results should contain, match exactly, start with, or end with the value you entered.
- Step 5** Click **Search**. Messages matching your search criteria are displayed below the Search section of the page.
- 

#### Related Topics

- [Searching Very Large Message Collections, page 31-23](#)

## Searching Very Large Message Collections

If you have a very large collection of messages in the spam quarantine, and if your search terms are not narrowly defined, your query may take a very long time to return information, or it may time out.

You will be prompted to confirm whether you want to resubmit your search. Please note that having multiple large searches running simultaneously can impact performance.

## Viewing Messages in the Spam Quarantine

The message listing shows messages in the spam quarantine. You can select how many messages are shown at one time. You can sort the display by clicking on the column headings. Click the same column again to reverse the sorting.

Click the subject of a message to view the message, including the body and headers. The message is displayed in the Message Details page. The first 20K of the message is displayed. If the message is longer, it is truncated at 20K and you can download the message via the link at the bottom of the message.

From the Message Details page you can delete a message (select **Delete**) or select **Release** to release the message. Releasing a message causes it to be delivered.

To view additional details about the message, click the **Message Tracking** link.

Note the following:

- **Viewing Messages with Attachments**

When viewing a message that includes an attachment, the body of the message is displayed, followed by a list of attachments.

- **Viewing HTML Messages**

The spam quarantine attempts to render an approximation of HTML-based messages. Images are not displayed.

- **Viewing Encoded Messages**

Base64-encoded messages are decoded and then displayed.

## Delivering Messages in the Spam Quarantine

To release a message for delivery, click the checkbox next to the message or messages that you want to release and select **Release** from the drop-down menu. Then click **Submit**.

Click the checkbox in the heading row to automatically select all messages currently displayed on the page.

Released messages proceed directly to the destination queue, skipping any further work queue processing in the email pipeline.

## Deleting Messages from the Spam Quarantine

The spam quarantine can be configured to automatically delete messages after a certain amount of time. Also, the spam quarantine can be configured to automatically delete the oldest messages once the quarantine has reached its maximum size. You may also delete messages from the spam quarantine manually.

To delete specific messages, click the checkbox next to the messages that you want to delete and then select **Delete** from the drop-down menu. Then click **Submit**. Click the checkbox in the heading row to automatically select all of the messages currently displayed on the page.

To delete all messages in the spam quarantine, disable the quarantine (see [About Disabling the Spam Quarantine, page 31-24](#)) and then click the **Delete All Messages** link. The number in parenthesis at the end of the link is the number of messages in the spam quarantine.

## Disk Space for the Spam Quarantine

By default, messages in the spam quarantine are automatically deleted after a set amount of time. If the quarantine gets full, older spam is deleted. To change this setting, see [Enabling and Configuring the Spam Quarantine, page 31-3](#).

### Related Topics

- [Managing Disk Space, page 33-16](#)

## About Disabling the Spam Quarantine

If you disable the spam quarantine:

- If messages are present in the spam quarantine when it is disabled, you can opt to delete all of the messages.
- Any mail policies set to quarantine spam or suspected spam will instead be set to deliver the message. You may need to adjust mail policies.
- To completely disable an external spam quarantine, disable it on both the Email Security appliance and the Security Management appliance.

Disabling an external spam quarantine on the Email Security appliance only does not delete the external quarantine or its messages and data.

## Troubleshooting Spam Quarantine Features

- [Troubleshooting Safelists and Blocklists, page 31-13](#)
- [Troubleshooting Spam Notifications, page 31-21](#)

- [Ensuring That Message Text Displays Correctly, page 31-6](#)





## Distributing Administrative Tasks

---

- [Working with User Accounts, page 32-1](#)
- [Managing Custom User Roles for Delegated Administration, page 32-7](#)
- [Passwords, page 32-16](#)
- [Configuring Access to the Email Security Appliance, page 32-24](#)
- [Managing Secure Shell \(SSH\) Keys, page 32-28](#)
- [Viewing Active Administrator Sessions, page 32-30](#)

### Working with User Accounts

The Cisco appliance provides two methods for adding user accounts: creating user accounts on the Cisco appliances itself, and enabling user authentication using your own centralized authentication system, which can be either an LDAP or RADIUS directory. You can manage users and connections to external authentication sources on the System Administration > Users page in the GUI (or by using the `userconfig` command in the CLI). For information about using an external directory to authenticate users, see [External Authentication, page 32-20](#).

The default user account for the system, admin, has all administrative privileges. The admin user account cannot be deleted, but you can change the password and lock the account.

When you create a new user account, you assign the user to a predefined or a custom user role. Each role contains differing levels of permissions within the system.

Although there is no limit to the number of user accounts that you can create on the appliance, you cannot create user accounts with names that are reserved by the system. For example, you cannot create the user accounts named “operator” or “root.”



## User Roles

**Table 32-1**      *User Roles Listing*

| User Role     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| admin         | <p>The <code>admin</code> user is the default user account for the system and has all administrative privileges. The <code>admin</code> user account is listed here for convenience, but it cannot be assigned via a user role, and it cannot be edited or deleted, aside from changing the password.</p> <p>Only the <code>admin</code> user can issue the <code>resetconfig</code> and <code>revert</code> commands.</p>                                                                                                                                                                                                                                                                                                                            |
| Administrator | <p>User accounts with the Administrator role have full access to all configuration settings of the system. However, only the <code>admin</code> user has access to the <code>resetconfig</code> and <code>revert</code> commands.</p> <p><b>Note</b> AsyncOS does not support multiple administrators configuring the Email Security appliance from the GUI simultaneously.</p>                                                                                                                                                                                                                                                                                                                                                                       |
| Technician    | <p>User accounts with the Technician role can perform system upgrades, reboot the appliance, and manage feature keys. Technicians can also perform the following actions in order to upgrade the appliance:</p> <ul style="list-style-type: none"> <li>• Suspend email delivery and receiving.</li> <li>• View status of workqueue and listeners.</li> <li>• Save and email configuration files.</li> <li>• Back up safelists and blocklists. Technicians cannot restore these lists.</li> <li>• Disconnect the appliance from a cluster.</li> <li>• Enable or disable remote service access for Cisco technical support.</li> <li>• Raise a support request.</li> </ul>                                                                              |
| Operator      | <p>User accounts with the Operator role are restricted from:</p> <ul style="list-style-type: none"> <li>• Creating or editing user accounts.</li> <li>• Issuing the <code>resetconfig</code> command.</li> <li>• Upgrading the appliance.</li> <li>• Issuing the <code>systemsetup</code> command or running the System Setup Wizard.</li> <li>• Issuing the <code>adminaccessconfig</code> command.</li> <li>• Performing some quarantine functions (including creating, editing, deleting, and centralizing quarantines).</li> <li>• Modifying LDAP server profile settings other than username and password, if LDAP is enabled for external authentication.</li> </ul> <p>Otherwise, they have the same privileges as the Administrator role.</p> |
| Guest         | <p>Users accounts with the Guest role can only view status information and reports. Users with the Guest role can also manage messages in quarantines, if access is enabled in a quarantine. Users with the Guest role cannot access Message Tracking.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

**Table 32-1**      **User Roles Listing**

| User Role          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Read-Only Operator | <p>User accounts with the Read-Only Operator role have access to view configuration information. Users with the Read-Only Operator role can make and submit changes to see how to configure a feature, but they cannot commit them. Users with this role can manage messages in quarantines, if access is enabled in a quarantine.</p> <p>Users with this role cannot access the following:</p> <ul style="list-style-type: none"> <li>• File system, FTP, or SCP.</li> <li>• Settings for creating, editing, deleting, or centralizing quarantines.</li> </ul>                                                            |
| Help Desk User     | <p>User accounts with the Help Desk User role are restricted to:</p> <ul style="list-style-type: none"> <li>• Message tracking.</li> <li>• Managing messages in quarantines.</li> </ul> <p>Users with this role cannot access to the rest of the system, including the CLI. You need to enable access in each quarantine before a user with this role can manage them.</p>                                                                                                                                                                                                                                                 |
| Custom user role   | <p>User accounts with a custom user role can only access email security features assigned to the role. These features can be any combination of DLP policies, email policies, reports, quarantines, local message tracking, encryption profiles, and the Trace debugging tool. The users cannot access system configuration features, including enabling features globally. Only administrators can define custom user roles. See <a href="#">Managing Custom User Roles for Delegated Administration, page 32-7</a> for more information.</p> <p><b>Note</b>    Users assigned to custom roles cannot access the CLI.</p> |

All roles defined in [Table 32-1](#) can access both the GUI and the CLI, except the Help Desk User role and custom user roles, which can only access the GUI.

If you use an LDAP directory to authenticate users, you assign directory groups to user roles instead of individual users. When you assign a directory group to a user role, each user in that group receives the permissions defined for the user role. For more information, see [External Authentication, page 32-20](#).

#### Related Topics

- [Managing Users, page 32-3](#)

## Managing Users

The Users page lists the existing users for the system, including the username, full name, and user type or group.

From the Users page, you can:

- Add new users. For more information, see [Adding Users, page 32-4](#).
- Delete users. For more information, see [Deleting Users, page 32-5](#).
- Edit users, such as changing a user's password and locking and unlocking a user's account. For more information, see [Editing Users, page 32-4](#).

- Force users to change their passwords. See [Force Users To Change Their Passwords, page 32-5](#).
- Configure user account and password settings for local accounts. For more information, see [Configuring Restrictive User Account and Password Settings, page 32-17](#).
- Enable the appliance to use an LDAP or RADIUS directory to authenticate users. For more information, see [External Authentication, page 32-20](#) for more information.
- Enable access for non-administrators to DLP Matched Content in Message Tracking. See [Controlling Access to Sensitive Information in Message Tracking, page 32-5](#) for more information.

#### Related Topics

- [Additional Commands to Support Multiple Users: who, whoami, and last, page 32-6](#)

## Adding Users

#### Before You Begin

- Determine the user roles you will use.
  - For descriptions of predefined user roles, see [User Roles, page 32-2](#).
  - To create custom roles, see [Managing Custom User Roles for Delegated Administration, page 32-7](#).
- Specify your password requirements. See [Configuring Restrictive User Account and Password Settings, page 32-17](#).

#### Procedure

- 
- Step 1** Choose **System Administration > Users**.
  - Step 2** Click **Add User**.
  - Step 3** Enter a login name for the user. Some words are reserved (such as “operator” or “root”).
  - Step 4** Enter the user’s full name.
  - Step 5** Select a predefined or custom user role.
  - Step 6** Generate or enter a password.
  - Step 7** Submit and commit your changes.
- 

## Editing Users

Use this procedure to change a password, etc.

#### Procedure

- 
- Step 1** Choose **System Administration > Users**.
  - Step 2** Click the user’s name in the Users listing.
  - Step 3** Make changes to the user.

- Step 4** Submit and commit your changes.
- 

## Force Users To Change Their Passwords

### Procedure

---

- Step 1** Choose **System Administration > Users**.
- Step 2** Select the users from the Users listing.
- Step 3** Click **Enforce Password Change**.
- Step 4** Choose whether the users must change the password during the next login or after a specified duration (in days).
- Step 5** (Optional) If you are enforcing a password change after a specified duration, set the grace period (in days) to reset the password after the password expires.
- Step 6** Click **OK**.
- Step 7** Submit and commit your changes.
- 

## Deleting Users

### Procedure

---

- Step 1** Click the trash can icon corresponding to the user's name in the Users listing.
- Step 2** Confirm the deletion by clicking **Delete** in the warning dialog that appears.
- Step 3** Commit your changes.
- 

## Controlling Access to Sensitive Information in Message Tracking

Messages that violate Data Loss Prevention (DLP) policies typically include sensitive information, such as corporate confidential information or personal information including credit card numbers and health records. By default, this content appears in the DLP Matched Content tab on the Message Details page for messages listed in Message Tracking results.

Administrator users can always see this content. However, you can choose to hide this tab and its content from users who have access to Message Tracking based on their assigned predefined or custom role.

### Before You Begin

Determine whether you have enabled matched content logging, which determines whether or not sensitive DLP data appears in Message Tracking. See [Showing or Hiding Sensitive DLP Data in Message Tracking](#), page 17-38.

**Procedure**

- 
- Step 1** Go to the **System Administration > Users** page.
- Step 2** Under **DLP Tracking Privileges**, click **Edit Settings**.
- Step 3** Select the roles for which you want to grant access to DLP data in Message Tracking.  
Custom roles without access to Message Tracking can never view this information and thus are not listed.
- Step 4** Submit and commit your changes.  
The following features must be enabled in Security Services for this setting to take effect:
- Message Tracking
  - RSA Email DLP
  - RSA Email DLP > Matched Content Logging
- 

**Additional Commands to Support Multiple Users: who, whoami, and last**

The following commands support multiple user access to the appliance.

- The `who` command lists all users who are logged into the system via the CLI, the time of login, the idle time, and the remote host from which the user is logged in:

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

```
Username Login Time Idle Time Remote Host What
===== ===== ===== ===== =====
admin 03:27PM 0s 10.1.3.201 cli
```

- The `whoami` command displays the username and full name of the user currently logged in, and which groups the user belongs to:

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

```
Username: admin
```

```
Full Name: Administrator
```

```
Groups: admin, operators, config, log, guest
```

- The `last` command displays which users have recently logged into the appliance. The IP address of the remote host, and the login, logout, and total time are also displayed.

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

| Username | Remote Host | Login Time       | Logout Time      | Total Time |
|----------|-------------|------------------|------------------|------------|
| =====    | =====       | =====            | =====            | =====      |
| admin    | 10.1.3.67   | Sat May 15 23:42 | still logged in  | 15m        |
| admin    | 10.1.3.67   | Sat May 15 22:52 | Sat May 15 23:42 | 50m        |
| admin    | 10.1.3.67   | Sat May 15 11:02 | Sat May 15 14:14 | 3h 12m     |
| admin    | 10.1.3.67   | Fri May 14 16:29 | Fri May 14 17:43 | 1h 13m     |
| shutdown |             |                  | Fri May 14 16:22 |            |
| shutdown |             |                  | Fri May 14 16:15 |            |
| admin    | 10.1.3.67   | Fri May 14 16:05 | Fri May 14 16:15 | 9m         |
| admin    | 10.1.3.103  | Fri May 14 16:12 | Fri May 14 16:15 | 2m         |
| admin    | 10.1.3.103  | Thu May 13 09:31 | Fri May 14 14:11 | 1d 4h 39m  |
| admin    | 10.1.3.135  | Fri May 14 10:57 | Fri May 14 10:58 | 0m         |
| admin    | 10.1.3.67   | Thu May 13 17:00 | Thu May 13 19:24 | 2h 24m     |

## Managing Custom User Roles for Delegated Administration

You can design custom user roles and delegate specific responsibilities to users that align with their roles within your organization, allowing these *delegated administrators* access only to the email security features they are responsible for and not the system configuration features that are not related to their roles. Delegated administration provides more flexible control over your users' access to the email security features on the appliance than the predefined administrator, operator, and help desk user roles.

For example, you may have users who are responsible for managing mail policies for specific domains on the Email Security appliance, but you do not want these users to access the system administration and security services configuration features, which the predefined administrator and operator roles grant. You can create a custom user role for mail policy administrators who can grant these users access to the mail policies they manage, along with other email security features that they can use to manage messages processed by these policies, such as Message Tracking and policy quarantines.

Use the System Administration > User Roles page in the GUI (or the `userconfig -> role` command in the CLI) to define custom user roles and manage the email security features for which they are responsible, such as mail policies, RSA Email DLP policies, email reports, and quarantines. For a full list of email security features that delegated administrators can manage, see [Assigning Access Privileges, page 32-9](#). Custom roles can also be created when adding or editing a local user account using the System Administration > Users page. See [Defining a Custom User Role When Adding a User Account, page 32-14](#) for more information.

You should make sure when creating a custom user role so that its responsibilities don't overlap too much with the responsibilities of other delegated administrators. If multiple delegated administrators are responsible for the same content filter, for example, and use the content filter in different mail policies, the changes made to the filter by one delegated administrator may cause unintended side effects for the mail policies managed by other delegated administrators.

When you have created the custom user roles, you can assign local users and external authentication groups to them like any other user role. See [Working with User Accounts, page 32-1](#) for more information. Please note that users assigned to custom roles cannot access the CLI.

[Figure 32-1](#) displays a list of custom user roles defined for an Email Security appliance, including the access privileges assigned to the roles.

**Figure 32-1** List of Custom User Roles  
User Roles

| Custom User Roles for Delegated Administration |                                                                                                                                                                                 |                                                                                                      |                   |                  |           |                |                     |                |                                                                                       |                                                                                       |
|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|-------------------|------------------|-----------|----------------|---------------------|----------------|---------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| Add User Role...                               |                                                                                                                                                                                 |                                                                                                      |                   |                  |           |                |                     |                |                                                                                       |                                                                                       |
| Role Name                                      | Privileges                                                                                                                                                                      |                                                                                                      |                   |                  |           |                |                     | Assigned Users | Duplicate                                                                             | Delete                                                                                |
|                                                | Email Policies                                                                                                                                                                  | Data Loss Prevention                                                                                 | Reporting         | Message Tracking | Trace     | Quarantines    | Encryption Profiles |                |                                                                                       |                                                                                       |
| DLP Administrator                              | No Access                                                                                                                                                                       | DLP Policies: 3<br> | Relevant Reports* | Available        | No Access | No Access      | Feature Disabled    | susan1         |    |    |
| Policy Administrator                           | Incoming Policies: 1<br>Content Filters: 0<br>Outgoing Policies: 1<br>Content Filters: 0<br> | No Access                                                                                            | Relevant Reports* | Available        | No Access | Quarantines: 1 | Feature Disabled    | grace1         |    |    |
| Quarantine Manager                             | No Access                                                                                                                                                                       | No Access                                                                                            | No Access         | No Access        | No Access | Quarantines: 3 | Feature Disabled    | jessie1        |  |  |

\* Report access for this role is controlled by the Mail Policy and DLP privileges.

Key:  View restricted to editable items

### Related Topics

- [Account Privileges Page, page 32-8](#)
- [Assigning Access Privileges, page 32-9](#)
- [Defining a Custom User Role, page 32-13](#)
- [Defining a Custom User Role When Adding a User Account, page 32-14](#)
- [Updating Responsibilities for a Custom User Role, page 32-14](#)
- [Editing a Custom User Role, page 32-15](#)
- [Duplicating a Custom User Role, page 32-15](#)
- [Deleting a Custom User Role, page 32-16](#)

## Account Privileges Page

When a delegated administrator logs into the appliance, the Account Privileges page displays links to the security features for which the delegated administrator is responsible and brief descriptions of their access privileges. A delegated administrator can return to this page by selecting Account Privileges in the Options menu. Delegated administrators can also access the features that they manage using the menu at the top of the web page.

Figure 32-2 shows an Account Privileges page for a delegated administrator with access to mail policies, email reporting, message tracking, and quarantines.

**Figure 32-2 Account Privileges Page for a Delegated Administrator**  
**Account Privileges (bob1)**

|                         |                                                                                                                                                                                                                                                                       |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Mail Policies</b>    | <a href="#">Incoming Mail Policies (1)</a><br><a href="#">Incoming Content Filters (1)</a><br><a href="#">Outgoing Mail Policies (1)</a><br><a href="#">Outgoing Content Filters (None Assigned)</a><br><a href="#">Configure Email Policies and Content Filters.</a> |
| <b>Email Reporting</b>  | <a href="#">Policy Reporting and DLP Reporting</a><br><a href="#">View and analyze email traffic.</a>                                                                                                                                                                 |
| <b>Message Tracking</b> | <a href="#">Message Tracking</a><br><a href="#">Track messages.</a>                                                                                                                                                                                                   |
| <b>Quarantine</b>       | <a href="#">Manage Message Quarantines (1)</a><br><a href="#">Manage messages in assigned Quarantines.</a>                                                                                                                                                            |

## Assigning Access Privileges

When creating a custom user role, you define the levels of access to the security features for which delegated administrators are responsible.

The security features available for delegated administrators to manage are:

- Incoming and outgoing mail policies and content filters.
- Data Loss Prevention (DLP) policies.
- Email reporting.
- Message Tracking.
- The Trace debugging tool.
- Spam, policy, virus, and outbreak quarantines.
- Cisco Email Encryption profiles.

After defining the access levels for a custom user role, you need to assign the specific mail policies, content filters, DLP policies, quarantines, or encryption profiles for which the delegated administrators will be responsible.

For example, you can create two different DLP policy administrator roles that are responsible for different RSA Email DLP policies. One role is only responsible for DLP violations related to company confidentiality and acceptable use, while the other is responsible for DLP violations related to privacy protection. In addition to DLP policies access, these custom user roles can also be assigned privileges for tracking message data and viewing quarantines and reports. They can search for DLP violations related to the policies that they are responsible for in using Message Tracking.

You can view which responsibilities are available to assign to a custom user role by clicking on the links for the assigned privileges in the Custom User Roles for Delegated Administration table on the User Roles page. See [Updating Responsibilities for a Custom User Role, page 32-14](#).



**Related Topics**

- [Mail Policies and Content Filters, page 32-10](#)
- [DLP Policies, page 32-11](#)
- [Email Reporting, page 32-12](#)
- [Message Tracking, page 32-12](#)
- [Trace, page 32-13](#)
- [Quarantines, page 32-13](#)
- [Encryption Profiles, page 32-13](#)

## Mail Policies and Content Filters

The Mail Policies and Content Filters access privileges define a delegated administrator's level of access to the incoming and outgoing mail policies and content filters on the Email Security appliance. You can assign specific mail policies and content filters to a custom user role, allowing only the delegated administrators belonging to this role, along with operators and administrators, to manage the mail policies and content filters.

All delegated administrators with this access privilege can view the default incoming and outgoing mail policies but they can only edit these policies if they have full access.

All delegated administrators with access privileges can create new content filters to use with their mail policies. A content filter created by a delegated administrator is available to the other delegated administrators assigned to the custom user role. Content filters that are not assigned to any custom user role are public and can be viewed by all delegated administrators with the mail policy access privilege. Content filters created by operators and administrators are *public* by default. Delegated administrators can enable or disable any existing content filters on mail policies assigned to their custom user role, but they cannot modify or delete public content filters.

If a delegated administrator deletes a content filter used by mail policies other than their own, or if the content filter is assigned to other custom user roles, AsyncOS does not delete the content filter from the system. AsyncOS instead unlinks the content filter from the custom user role and removes it from the delegated administrator's mail policies. The content filter remains available to other custom user roles and mail policies.

Delegated administrators can use any text resource or dictionary in their content filters, but they cannot access the Text Resources or Dictionaries pages in the GUI to view or modify them. Delegated administrators also cannot create new text resources or dictionaries.

For outgoing mail policies, delegated administrators can enable or disable DLP policies but they cannot customize the DLP settings unless they also have DLP policy privileges.

You can assign one of the following access levels for mail policies and content filters to a custom user role:

- **No access:** Delegated administrators cannot view or edit mail policies and content filters on the Email Security appliance.
- **View assigned, edit assigned:** Delegated administrators can view and edit the mail policies and content filters assigned to the custom user role and create new content filters. Delegated administrators can edit a policy's Anti-Spam, Anti-Virus, and Outbreak Filters settings. They can enable their content filters for the policy, as well as disable any existing content filter assigned to the policy, regardless of whether they are responsible for it. Delegated administrators cannot modify a mail policy's name or its senders, recipients, or groups. Delegated administrators can modify the order of the content filters for mail policies assigned to their custom user role.

- **View all, edit assigned:** Delegated administrators can view all mail policies and content filters on the appliance, but they can only edit the ones assigned to the custom user role.

**View all, edit all (full access):** Delegated administrators have full access to all of the mail policies and content filters on the appliance, including the default mail policies, and have the ability to create new mail policies. Delegated administrators can modify the senders, recipients, and groups of all mail policies. They can also reorder mail policies.

You can assign individual mail policies and content filters to the custom user role using either the Email Security Manager or the Custom User Roles for Delegated Administration table on the User Roles page.

See [Updating Responsibilities for a Custom User Role, page 32-14](#) for information on using the Custom User Roles for Delegated Administration table to assign mail policies and content filters.

## DLP Policies

The DLP Policies access privileges define a delegated administrator's level of access to the DLP policies via the DLP Policy Manager on the Email Security appliance. You can assign DLP policies to specific custom user roles, allowing delegated administrators, in addition to operators and administrators, to manage these policies. Delegated administrators with DLP access can also export DLP configuration files from the Data Loss Prevention Global Settings page. Only administrators and operators can change the mode of DLP used from RSA Email DLP to RSA Enterprise Manager, and vice versa.

If a delegated administrator also has mail policy privileges, they can customize the RSA Email DLP policies. Delegated administrators can use any custom DLP dictionary for their RSA Email DLP policies, but they cannot view or modify the custom DLP dictionaries.

You can assign one of the following access levels for RSA Email DLP policies to a custom user role:

- **No access:** Delegated administrators cannot view or edit RSA Email DLP policies on the Email Security appliance.
- **View assigned, edit assigned:** Delegated administrators can use the DLP Policy Manager to view and edit the RSA Email DLP policies assigned to the custom user role. Delegated administrators cannot rename or reorder DLP policies in the DLP Policy Manager. Delegated administrators can export DLP configurations.
- **View all, edit assigned:** Delegated administrators can view and edit the RSA Email DLP policies assigned to the custom user role. They can export DLP configurations. They can also view all RSA Email DLP policies that are not assigned to the custom user role but they cannot edit them. Delegated administrators cannot reorder DLP policies in the DLP Policy Manager or rename the policy.
- **View all, edit all (full access):** Delegated administrators have full access to all of the RSA Email DLP policies on the appliance, including the ability to create new ones. Delegated administrators can reorder DLP policies in the DLP Policy Manager. They cannot change the DLP mode that the appliance uses.

You can assign individual RSA Email DLP policies to the custom user role using either the DLP Policy Manager or the Custom User Roles for Delegated Administration table on the User Roles page.

See [Chapter 17, "Data Loss Prevention"](#) for more information on RSA Email DLP policies and the DLP Policy Manager.

See [Updating Responsibilities for a Custom User Role, page 32-14](#) for information on using the Custom User Roles for Delegated Administration list to assign RSA Email DLP policies.

## Email Reporting

The Email Reporting access privileges define which reports and Email Security Monitor pages a delegated administrator can view, depending on the custom user role's access to mail policies, content filters, and RSA Email DLP policies. These reports are not filtered for assigned policies; delegated administrators can view reports for mail and DLP policies that for which they are not responsible.

You can assign one of the following access levels for email reporting to a custom user role:

- **No access:** Delegated administrators cannot view reports on the Email Security appliance.
- **View relevant reports:** Delegated administrators can view reports on the Email Security Monitor pages related to their Mail Policies and Content Filters and DLP Policies access privileges. Delegated administrators with Mail Policies and Content Filters access privileges can view the following Email Security Monitor pages:
  - Overview
  - Incoming Mail
  - Outgoing Destinations
  - Outgoing Senders
  - Internal Users
  - Content Filters
  - Virus Outbreaks
  - Virus Types
  - Archived Reports

Delegated administrators with DLP Policies access privileges can view the following Email Security Monitor pages:

- Overview
- DLP Incidents
- Archived Reports
- **View all reports:** Delegated administrators can view all reports and Email Security Monitor pages on the Email Security appliance.

See the Chapter 28, “Using Email Security Monitor,” on page 1 chapter for more information on email reporting and the Email Security Monitor.

## Message Tracking

The Message Tracking access privileges define whether delegated administrators assigned to the custom user role have access to Message Tracking, including message content that may violate your organization's DLP policies if the DLP Tracking Policies option has been enabled on the System Administration > Users page and the custom user role also has DLP policies access privileges.

Delegated administrators can only search for the DLP violations for the RSA Email DLP policies assigned to them.

See Chapter 29, “Tracking Messages,” on page 1 for more information on Message Tracking.

See [Controlling Access to Sensitive Information in Message Tracking, page 32-5](#) for information for allowing delegated administrators access to viewing matched DLP content in Message Tracking.

## Trace

The Trace access privileges define whether delegated administrators assigned to the custom user role can use Trace to debug the flow of messages through the system. Delegated administrators with access can run Trace and view all of the generated output. Trace results are not filtered based on the delegated administrator's mail or DLP policy privileges.

See [Debugging Mail Flow Using Test Messages: Trace, page 40-1](#) for more information on using Trace.

## Quarantines

The Quarantines access privileges define whether delegated administrators can manage assigned quarantines. Delegated administrators can view and take actions on any message in an assigned quarantine, such as releasing or deleting messages, but cannot change the quarantine's configuration (e.g. the size, retention period, etc.), or create or delete quarantines.

You can assign any of the quarantines to the custom user role using either the Monitor > Quarantines page or the Custom User Roles for Delegated Administration table on the User Roles page.

See [About Distributing Message Processing Tasks to Other Users, page 30-9](#) and [Configuring Administrative User Access to the Spam Quarantine, page 31-4](#) for more information on assigning Quarantine management tasks to administrative users.

See [Updating Responsibilities for a Custom User Role, page 32-14](#) for information on using the Custom User Roles for Delegated Administration list to assign quarantines.

## Encryption Profiles

The Encryption Profiles access privileges define whether delegated administrators can use encryption profiles assigned to their custom user role when editing content filters or DLP policies. Encryption profiles can only be assigned to custom user roles with mail or DLP policy access privileges. Encryption profiles that are not assigned to a custom role are available for use by all delegated administrators with mail or DLP policy privileges. Delegated administrators cannot view or modify any encryption profiles.

You can assign encryption profiles when creating or editing an encryption profile using the Security Services > IronPort Email Encryption page.

## Defining a Custom User Role

User the User Roles page in the GUI (or the `userconfig -> role` command in the CLI) to define a new user role and assign its access privileges. The User Roles page displays all existing custom user roles on the appliance and the access privileges for each role.

### Procedure

- 
- Step 1** Choose **System Administration > User Roles**.
  - Step 2** Click **Add User Role**.
  - Step 3** Enter a name for the user role.
  - Step 4** Enter a description of the user role and its privileges.
  - Step 5** Select the user role's access privileges. (See [Assigning Access Privileges, page 32-9](#) for more information on each type of access privilege.)

- Step 6** Submit and commit your changes.
- 

## Defining a Custom User Role When Adding a User Account

You can create a new custom user role when adding or editing a local user account on the Email Security appliance.

See [Managing Users, page 32-3](#) for more information on adding a user account.

### Procedure

---

- Step 1** Go to the **System Administration > Users** page.
  - Step 2** Click **Add User**.
  - Step 3** When creating the user account, select Custom Roles.
  - Step 4** Select **Add Role**.
  - Step 5** Enter the name for the new role.
  - Step 6** Submit the new user account.  
AsyncOS displays a notification that the new user account and custom user role have been added.
  - Step 7** Go to the **System Administration > User Roles** page.
  - Step 8** Click on the name of the custom user role in the Custom User Roles for Delegated Administration table.
  - Step 9** Enter a description of the user role and its privileges.
  - Step 10** Select the user role's access privileges. (See [Assigning Access Privileges, page 32-9](#) for more information on each type of access privilege.)
  - Step 11** Submit and commit your changes.
- 

## Updating Responsibilities for a Custom User Role

While you can assign responsibilities to custom user roles by browsing to the individual security features using the menu at the top of the GUI, the Custom User Roles for Delegated Administration table on the User Roles page consolidates links to all of the security features that delegated administrators can manage in one place, with the exception of Encryption profiles. Clicking on the name of a custom user group's access privilege in the table displays a list of all the mail policies, content filters, active RSA Email DLP policies, or quarantines on the appliance and displays the names of any other custom user role that has access to them.

For example, [Figure 32-3](#) displays a list of active RSA Email DLP policies available on an Email Security appliance. It also lists another custom user group that has access to the DLP policies. From this list, an administrator can select which DLP policies the delegated administrators using the DLP Policy Manager.

**Figure 32-3 DLP Policies Available for Delegated Administrators**  
**User Role: DLP Administrator > DLP Policies**

| Active DLP Policies for Outgoing Mail |       |                                                        |                              |
|---------------------------------------|-------|--------------------------------------------------------|------------------------------|
| Include                               | Order | DLP Policy                                             | Other Roles with Edit Access |
| <input checked="" type="checkbox"/>   | 1     | Payment Card Industry Data Security Standard (PCI-DSS) | Domain Admin                 |
| <input checked="" type="checkbox"/>   | 2     | California SB-1386                                     | Domain Admin                 |
| <input type="checkbox"/>              | 3     | Restricted Files                                       | Domain Admin                 |

Cancel

Submit

### Procedure

- 
- Step 1** Go to the **System Administration > User Roles** page.
- Step 2** Click the name of the access privilege for the custom user role you want to update.
- AsyncOS displays a list of all the mail policies, content filters, DLP policies, or quarantines available on the appliance, along with the names of any other assigned custom user roles.
- Step 3** Select the mail policies, content filters, DLP policies, or quarantines for which you want the delegated administrators assigned to be responsible.
- Step 4** Submit and commit your changes.
- 

## Editing a Custom User Role

### Procedure

- 
- Step 1** Go to the **System Administration > User Roles** page.
- Step 2** Click the user role's name in the Custom User Roles for Delegated Administration listing.
- Step 3** Make changes to the user role.
- Step 4** Submit and commit your changes.
- 

## Duplicating a Custom User Role

You may want to create multiple custom user roles with similar access privileges but assign different responsibilities to different sets of users. For example, if the Email Security appliance handles messages for multiple domains, you may want to create custom user roles with similar access rights but for different mail policies based on the domain. This allows delegated administrators to manage mail policies for their domains without interfering with the responsibilities of other delegated administrators.

### Procedure

- 
- Step 1** Go to the **System Administration > User Roles** page.
- Step 2** Click the duplicate icon corresponding to the user role you want to duplicate in the Custom User Roles for Delegated Administration listing.

- Step 3** Change the name of the custom user role.
  - Step 4** Make any access privilege changes required for the new custom user role.
  - Step 5** Submit and commit your changes.
- 

## Deleting a Custom User Role

When a custom role is deleted, users become unassigned and do not have access to the appliance. If you delete a custom user role that is assigned to one or more users, you do not receive a warning message. You should reassign any users that were assigned to the custom user role that you deleted.

### Procedure

- Step 1** Go to the **System Administration > User Roles** page.
  - Step 2** Click the trash can icon corresponding to the user role you want to delete in the Custom User Roles for Delegated Administration list.
  - Step 3** Confirm the deletion by clicking **Delete** in the warning dialog that appears.
  - Step 4** Commit your changes.
- 

## Passwords

- [Changing Your Password, page 32-16](#)
- [Locking and Unlocking a User Account, page 32-17](#)
- [Configuring Restrictive User Account and Password Settings, page 32-17](#)
- [External Authentication, page 32-20](#)

## Changing Your Password

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

As soon as you submit a new password, you are logged out and taken to the log in screen.

In the CLI, use the `password` or `passwd` command to change your password. If you forget the password for the admin user account, contact your customer support provider to reset the password.

The `password` command requires you to enter the old password for security.



### Note

Changes to the password take effect immediately and do not require you commit the change.

---

## Locking and Unlocking a User Account

Locking a user account prevents a local user from logging into the appliance. A user account can be locked in one of the following ways:

- AsyncOS locks a user account if the user exceeded the maximum number of failed login attempts defined in the Local User Account & Password Settings section.
- Administrators can manually lock user accounts for security purposes using the System Administration > Users page.

AsyncOS displays the reason why the user account was locked when you view the user account on the Edit User page.

To unlock a user account, open the user account by clicking on the user name in the Users listing and click **Unlock Account**.

To manually lock a local user account, open the user account by clicking on the user name in the Users listing and click **Lock Account**. AsyncOS displays a message saying that the user will be unable to log into the appliance and asks if you want to continue.

You can also configure all local user accounts to lock after users fail to login successfully after a configured number of attempts. For more information, see [Configuring Restrictive User Account and Password Settings, page 32-17](#).

**Note**

If you lock the admin account, you can only unlock it by logging in as the admin through a serial communications connection to the serial console port. The admin user can always access the appliance using the serial console port, even when the admin account is locked. See [Connecting to the Appliance, page 3-9](#) for more information on accessing the appliance using the serial console port.

## Configuring Restrictive User Account and Password Settings

You can define user account and password restrictions to enforce organizational password policies. The user account and password restrictions apply to local users defined on the Cisco appliance. You can configure the following settings:

- **User account locking.** You can define how many failed login attempts cause the user to be locked out of the account.
- **Password lifetime rules.** You can define how long a password can exist before the user is required to change the password after logging in.
- **Password rules.** You can define what kinds of passwords users can choose, such as which characters are optional or mandatory.

You define user account and password restrictions on the System Administration > Users page in the Local User Account and Password Settings section.

**Procedure**

- 
- Step 1** Choose **System Administration > Users**.
  - Step 2** Scroll to the **Local User Account and Password Settings** section.
  - Step 3** Click **Edit Settings**.



**Step 4** Configure the settings as described below.

| Setting           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| User Account Lock | <p>Choose whether or not to lock the user account after the user fails to login successfully. Specify the number of failed login attempts that cause the account locking. You can enter any number from one (1) to 60. Default is five (5).</p> <p>When you configure account locking, enter the message to be displayed to the user attempting to login. Enter text using 7-bit ASCII characters. This message is only displayed when users enter the correct password to an account locked by an administrator. This message is not shown for accounts locked due to failed login attempts.</p> <p>When a user account gets locked, an administrator can unlock it on the Edit User page in the GUI or using the <code>userconfig</code> CLI command.</p> <p>Failed login attempts are tracked by user, regardless of the machine the user connects from or the type of connection, such as SSH or HTTP. Once the user successfully logs in, the number of failed login attempts is reset to zero (0).</p> <p>When a user account is locked out due to reaching the maximum number of failed login attempts, an alert is sent to the administrator. The alert is set at the “Info” severity level.</p> <p><b>Note</b> You can also manually lock individual user accounts. For more information see <a href="#">Locking and Unlocking a User Account, page 32-17</a>.</p>                                                                       |
| Password Reset    | <p>You can choose whether:</p> <ul style="list-style-type: none"> <li>• Users should be forced to change their passwords after an administrator changes their passwords.</li> <li>• Users should be forced to change their passwords after a specified duration. Enter the number of days a password can last before users must change it. You can enter any number from one (1) to 366. Default is 90. In this case, you can optionally choose: <ul style="list-style-type: none"> <li>– To display a notification about the upcoming password expiration. Enter the number of days before expiration to notify users.</li> <li>– To allow a grace period (of specified days) to reset the password after the password expiry. Enter the number of days.</li> </ul> <p>If you are setting a grace period, user accounts will be locked if the passwords are not changed within the specified duration. If you are not setting a grace period, users can change their passwords any time after the password expiry.</p> </li> </ul> <p><b>Note</b> When a user account uses SSH keys instead of a password challenge, the Password Reset rules still apply. When a user account with SSH keys expires, the user must enter their old password or ask an administrator to manually change the password to change the keys associated with the account. For more information, see <a href="#">Managing Secure Shell (SSH) Keys, page 32-28</a>.</p> |

| Setting                                                             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Password Rules:<br>Require at least <number> characters.            | Enter the minimum number of characters that a password may contain.<br>Enter any number between 0 and 128.<br>Default is 8 characters.<br>Passwords can have more characters than the number you specify here.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Password Rules:<br>Require at least one number (0-9).               | Choose whether or not the passwords must contain at least one number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Password Rules:<br>Require at least one special character.          | Choose whether or not the passwords must contain at least one special character. Passwords may contain the following special characters:<br>~ ? ! @ # \$ % ^ & * - _ + =<br>\   / [ ] ( ) < > { } ` ' " ; : , .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Password Rules:<br>Ban usernames and their variations as passwords. | Choose whether or not the password are allowed to be the same as the associated username or variations on the username. When username variations are banned, the following rules apply to passwords: <ul style="list-style-type: none"> <li>• The password may not be the same as the username, regardless of case.</li> <li>• The password may not be the same as the username in reverse, regardless of case.</li> <li>• The password may not be the same as the username or reversed username with the following character substitutions:               <ul style="list-style-type: none"> <li>– "@" or "4" for "a"</li> <li>– "3" for "e"</li> <li>– "l", "!", or "1" for "i"</li> <li>– "0" for "o"</li> <li>– "\$" or "5" for "s"</li> <li>– "+" or "7" for "t"</li> </ul> </li> </ul> |
| Password Rules:<br>Ban reuse of the last <number> passwords.        | Choose whether or not users are allowed to choose a recently used password when they are forced to change the password. If they are not allowed to reuse recent passwords, enter the number of recent passwords that are banned from reuse.<br>You can enter any number from one (1) to 15. Default is three (3).                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

| Setting                                                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Password Rules:<br><br>List of words to disallow in passwords | <p>You can create a list of words to disallow in passwords.</p> <p>Make this file a text file with each forbidden word on a separate line. Save the file with the name <code>forbidden_password_words.txt</code> and use SCP or FTP to upload the file to the appliance.</p> <p>If this restriction is selected but no word list is uploaded, this restriction is ignored.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Password Strength                                             | <p>You can display a password-strength indicator when an admin or user enters a new password.</p> <p>This setting does not enforce creation of strong passwords, it merely shows how easy it is to guess the entered password.</p> <p>Select the roles for which you wish to display the indicator. Then, for each selected role, enter a number greater than zero. A larger number means that a password that registers as strong is more difficult to achieve. This setting has no maximum value.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• If you enter 30, then an 8 character password with at least one upper- and lower-case letter, number, and special character will register as a strong password.</li> <li>• If you enter 18, then an 8 character password with all lower case letters and no numbers or special characters will register as strong.</li> </ul> <p>Password strength is measured on a logarithmic scale. Evaluation is based on the U.S. National Institute of Standards and Technology rules of entropy as defined in NIST SP 800-63, Appendix A.</p> <p>Generally, stronger passwords:</p> <ul style="list-style-type: none"> <li>• Are longer</li> <li>• Include upper case, lower case, numeric, and special characters</li> <li>• Do not include words in any dictionary in any language.</li> </ul> <p>To enforce passwords with these characteristics, use the other settings on this page.</p> |

**Step 5** Submit and commit your changes.

### What To Do Next

If you selected **List of words to disallow in passwords**, create and upload the described text file.

## External Authentication

If you store user information in an LDAP or RADIUS directory on your network, you can configure your Cisco appliance to use the external directory to authenticate users who log in to the appliance. To set up the appliance to use an external directory for authentication, use the System Administration > Users page in the GUI or the `userconfig` command and the `external` subcommand in the CLI.

When external authentication is enabled and a user logs into the Email Security appliance, the appliance first determines if the user is the system defined “admin” account. If not, then the appliance checks the first configured external server to determine if the user is defined there. If the appliance cannot connect to the first external server, the appliance checks the next external server in the list.

For LDAP servers, if the user fails authentication on any external server, the appliance tries to authenticate the user as a local user defined on the Email Security appliance. If the user does not exist on any external server or on the appliance, or if the user enters the wrong password, access to the appliance is denied.

If an external RADIUS server cannot be contacted, the next server in the list is tried. If all servers cannot be contacted, the appliance tries to authenticate the user as a local user defined on the Email Security appliance. However, if an external RADIUS server rejects a user for any reason, such as an incorrect password or the user being absent, access to the appliance is denied.

#### Related Topics

- [Enabling LDAP Authentication, page 32-21](#)
- [Enabling RADIUS Authentication, page 32-22](#)

## Enabling LDAP Authentication

In addition to using an LDAP directory to authenticate users, you can assign LDAP groups to Cisco user roles. For example, you can assign users in the IT group to the Administrator user role, and you can assign users in the Support group to the Help Desk User role. If a user belongs to multiple LDAP groups with different user roles, AsyncOS grants the user the permissions for the most restrictive role. For example, if a user belongs to a group with Operator permissions and a group with Help Desk User permissions, AsyncOS grants the user the permissions for the Help Desk User role.



#### Note

If an external user changes the user role for their LDAP group, the user should log out of the appliance and then log back in. The user will have the permissions of their new role.

#### Before You Begin

Define an LDAP server profile and an external authentication query for the LDAP server. For more information, see [Chapter 25, “LDAP Queries.”](#)

#### Procedure

- Step 1** Choose **System Administration > Users**.
- Step 2** Scroll down to the **External Authentication** section.
- Step 3** Click **Enable**.
- Step 4** Select the **Enable External Authentication** check box.
- Step 5** Select **LDAP** for the authentication type.
- Step 6** Enter the amount of time to store external authentication credentials in the web user interface.
- Step 7** Select the LDAP external authentication query that authenticates users.
- Step 8** Enter the number of seconds that the appliance waits for a response from the server before timing out.
- Step 9** Enter the name of a group from the LDAP directory that you want the appliance to authenticate, and select the role for the users in the group.

- Step 10** Optionally, click **Add Row** to add another directory group. Repeat steps 9 and 10 for each directory group that the appliance authenticates.
- Step 11** Submit and commit your changes.

## Enabling RADIUS Authentication

You can also use a RADIUS directory to authenticate users and assign groups of users to Cisco roles. The RADIUS server should support the CLASS attribute, which AsyncOS uses to assign users in the RADIUS directory to Cisco user roles. AsyncOS supports two authentication protocols for communicating with the RADIUS server: Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP).

To assign RADIUS users to Cisco user roles, first set the CLASS attribute on the RADIUS server with a string value of <radius-group>, which will be mapped to Cisco user roles. The CLASS attribute may contain letters, numbers, and a dash, but cannot start with a dash. AsyncOS does not support multiple values in the CLASS attribute. RADIUS users belonging to a group without a CLASS attribute or an unmapped CLASS attribute cannot log into the appliance.

If the appliance cannot communicate with the RADIUS server, the user can log in with a local user account on the appliance.



### Note

If an external user changes the user role for their RADIUS group, the user should log out of the appliance and then log back in. The user will have the permissions of their new role.

### Procedure

- Step 1** On the System Administration > Users page, click **Enable**.
- Step 2** Check the **Enable External Authentication** option if it is not enabled already.
- Step 3** Enter the hostname for the RADIUS server.
- Step 4** Enter the port number for the RADIUS server. The default port number is 1812.
- Step 5** Enter the Shared Secret password for the RADIUS server.
- Step 6** Enter the number of seconds for the appliance to wait for a response from the server before timing out.
- Step 7** (Optional) Click **Add Row** to add another RADIUS server. Repeat steps 3–6 for each RADIUS server.



### Note

You can add up to ten RADIUS servers.

- Step 8** Enter the number of seconds AsyncOS stores the external authentication credentials before contacting the RADIUS server again to re-authenticate in the “External Authentication Cache Timeout” field. Default is zero (0).



### Note

If the RADIUS server uses one-time passwords, for example passwords created from a token, enter zero (0). When the value is set to zero, AsyncOS does not contact the RADIUS server again to authenticate during the current session.

**Step 9** Configure Group Mapping:

| Setting                                                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Map externally authenticated users to multiple local roles.       | <p>AsyncOS assigns RADIUS users to appliance roles based on the RADIUS CLASS attribute. CLASS attribute requirements:</p> <ul style="list-style-type: none"> <li>• 3 character minimum</li> <li>• 253 character maximum</li> <li>• no colons, commas, or newline characters</li> <li>• one or more mapped CLASS attributes for each RADIUS user (With this setting, AsyncOS denies access to RADIUS users without a mapped CLASS attribute.)</li> </ul> <p>For RADIUS users with multiple CLASS attributes, AsyncOS assigns the most restrictive role. For example, if a RADIUS user has two CLASS attributes, which are mapped to the Operator and Read-Only Operator roles, AsyncOS assigns the RADIUS user to the Read-Only Operator role, which is more restrictive than the Operator role.</p> <p>These are the appliance roles ordered from least restrictive to most restrictive:</p> <ul style="list-style-type: none"> <li>• admin</li> <li>• Administrator</li> <li>• Technician</li> <li>• Operator</li> <li>• Read-only Operator</li> <li>• Help Desk User</li> <li>• Guest</li> </ul> |
| Map all externally authenticated users to the Administrator role. | AsyncOS assigns RADIUS users to the Administrator role.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

**Step 10** Choose whether to map all externally authenticated users to the Administrator role or to different appliance user role types.

**Step 11** If you map users to different role types, enter the group name as defined in the RADIUS CLASS attribute in the Group Name or Directory field, and choose an appliance role type from the Role field. You can add more role mappings by clicking **Add Row**.

For more information on user role types, see [Working with User Accounts, page 32-1](#).

**Step 12** Submit and commit your changes.

# Configuring Access to the Email Security Appliance

AsyncOS provides administrators controls to manage users' access to the Email Security appliance, including a timeout for Web UI session and an access list that specifies the IP addresses from which users and your organization's proxy servers can access the appliance.

## Related Topics

- [Configuring IP-Based Network Access, page 32-24](#)
- [Configuring Session Timeouts, page 32-26](#)
- [Displaying Messages to Administrative Users, page 32-27](#)
- [Displaying a Message After Login, page 32-27](#)

## Configuring IP-Based Network Access

You can control from which IP addresses users access the Email Security appliance by creating access lists for users who connect directly to the appliance and users who connect through a reverse proxy, if your organization uses reverse proxies for remote users.

## Related Topics

- [Direct Connections, page 32-24](#)
- [Connecting Through a Proxy, page 32-24](#)
- [Creating the Access List, page 32-25](#)

## Direct Connections

You can specify the IP addresses, subnets, or CIDR addresses for machines that can connect to the Email Security appliance. Users can access the appliance from any machine with IP address from the access list. Users attempting to connect to the appliance from an address not included in the list are denied access.

## Connecting Through a Proxy

If your organization's network uses reverse proxy servers between remote users' machines and the Email Security appliance, AsyncOS allows you create an access list with the IP addresses of the proxies that can connect to the appliance.

Even when using a reverse proxy, AsyncOS still validates the IP address of the remote user's machine against a list of IP addresses allowed for user connections. To send the remote user's IP address to the Email Security appliance, the proxy needs to include the `x-forwarded-for` HTTP header in its connection request to the appliance.

The `x-forwarded-for` header is a non-RFC standard HTTP header with the following format:

```
x-forwarded-for: client-ip, proxy1, proxy2,... CRLF.
```

The value for this header is a comma-separated list of IP addresses with the left-most address being the address of the remote user's machine, followed by the addresses of each successive proxy that forwarded the connection request. (The header name is configurable.) The Email Security appliance matches the remote user's IP address from the header and the connecting proxy's IP address against the allowed user and proxy IP addresses in the access list.

**Note**

AsyncOS supports only IPv4 addresses in the `x-forwarded-for` header.

## Creating the Access List

You can create the network access list either via the Network Access page in the GUI or the `adminaccessconfig > ipaccess` CLI command.

AsyncOS offers four different modes of control for the access list:

- **Allow All.** This mode allows all connections to the appliance. This is the default mode of operation.
- **Only Allow Specific Connections.** This mode allows a user to connection to the appliance if the user's IP address matches the IP addresses, IP ranges, or CIDR ranges included in the access list.
- **Only Allow Specific Connections Through Proxy.** This mode allows a user to connect to the appliance through a reverse proxy if the following conditions are met:
  - The connecting proxy's IP address is included in the access list's IP Address of Proxy Server field.
  - The proxy includes the `x-forwarded-header` HTTP header in its connection request.
  - The value of `x-forwarded-header` is not empty.
  - The remote user's IP address is included in `x-forwarded-header` and it matches the IP addresses, IP ranges, or CIDR ranges defined for users in the access list.
- **Only Allow Specific Connections Directly or Through Proxy.** This mode allows users to connect through a reverse proxy or directly to the appliance if their IP address matches the IP addresses, IP ranges, or CIDR ranges included in the access list. The conditions for connecting through a proxy are the same as in the Only Allow Specific Connections Through Proxy mode.

Please be aware that you may lose access to the appliance after submitting and committing your changes if one of the following conditions is true:

- If you select **Only Allow Specific Connections** and do not include the IP address of your current machine in the list.
- If you select **Only Allow Specific Connections Through Proxy** and the IP address of the proxy currently connected to the appliance is not in the proxy list and the value of the Origin IP header is not in the list of allowed IP addresses.
- If you select **Only Allow Specific Connections Directly or Through Proxy** and
  - the value of the Origin IP header is not in the list of allowed IP addressesOR
  - the value of the Origin IP header is not in the list of allowed IP Addresses and the IP address of the proxy connected to the appliance is not in the list of allowed proxies.

### Procedure

**Step 1** Select **System Administration > Network Access**.

**Step 2** Click **Edit Settings**.

**Step 3** Select the mode of control for the access list.

**Step 4** Enter the IP addresses from which users will be allowed to connect to the appliance.

You can enter an IP address, IP address range or CIDR range. Use commas to separate multiple entries.



- Step 5** If connecting through a proxy is allowed, enter the following information:
- The IP addresses of the proxies allowed to connect to the appliance. Use commas to separate multiple entries.
  - The name of the origin IP header that the proxy sends to the appliance, which contains the IP addresses of the remote user's machine and the proxy servers that forwarded the request. By default, the name of the header is `x-forwarded-for`.
- Step 6** Submit and commit your changes.
- 

## Configuring Session Timeouts

- [Configuring the Web UI Session Timeout, page 32-26](#)
- [Configuring the CLI Session Timeout, page 32-26](#)

### Configuring the Web UI Session Timeout

You can specify how long a user can be logged into the Email Security appliance's Web UI before AsyncOS logs the user out due to inactivity. This Web UI session timeout applies to:

- All users, including administrator
- HTTP and HTTPS sessions
- Cisco Spam Quarantine

Once AsyncOS logs a user out, the appliance redirects the user's web browser to login page.

#### Procedure

---

- Step 1** Select **System Administration > Network Access**.
- Step 2** Click **Edit Settings**.
- Step 3** In the **Web UI Inactivity Timeout** field, enter the number of minutes users can be inactive before being logged out. You can define a timeout period between 5 and 1440 minutes.
- Step 4** Submit and commit your changes.
- 

You can also use the `adminaccessconfig` command in CLI to configure Web UI session timeout. See *Cisco AsyncOS for Email CLI Reference Guide*.

### Configuring the CLI Session Timeout

You can specify how long a user can be logged into the Email Security appliance's CLI before AsyncOS logs the user out due to inactivity. The CLI session timeout applies:

- To all users, including administrator
- Only to the connections using Secure Shell (SSH), SCP, and direct serial connection

**Note**

Any uncommitted configuration changes at the time of CLI session timeout will be lost. Make sure that you commit the configuration changes as soon as they are made.

**Procedure**

- 
- Step 1** Select **System Administration > Network Access**.
  - Step 2** Click **Edit Settings**.
  - Step 3** In the **CLI Inactivity Timeout** field, enter the number of minutes users can be inactive before being logged out. You can define a timeout period between 5 and 1440 minutes.
  - Step 4** Submit and commit your changes.
- 

You can also use the `adminaccessconfig` command in CLI to configure CLI session timeout. See *Cisco AsyncOS for Email CLI Reference Guide*.

## Displaying Messages to Administrative Users

### Displaying a Message Before Login

You can configure the Email Security appliance to display a message before a user attempts to log into the appliance through SSH, Telnet, FTP, or Web UI. The login banner is customizable text that appears above the login prompt. You can use the login banner to display internal security information or best practice instructions for the appliance. For example, you can create a simple note that saying that unauthorized use of the appliance is prohibited or a detailed warning concerning the organization's right to review changes made by the user to the appliance.

Use the `adminaccessconfig > banner` command in the CLI to create the login banner. The maximum length of the login banner is 2000 characters to fit 80x25 consoles. A login banner can be imported from a file in the `/data/pub/configuration` directory on the appliance. After creating the banner, commit your changes.

### Displaying a Message After Login

You can configure AsyncOS for Email to display a welcome banner after a user successfully logs into the appliance through SSH, Telnet, FTP, or Web UI. You can use the welcome banner to display internal security information or best practice instructions for the appliance.

Use the `adminaccessconfig > welcome` command in CLI to create the welcome banner. The maximum length of the welcome banner is 1600 characters.

You can import a welcome banner from a file in the `/data/pub/configuration` directory on your appliance. After creating the banner, commit your changes.

For more information, see *Cisco AsyncOS for Email CLI Reference Guide*.

# Managing Secure Shell (SSH) Keys

Use the `sshconfig` command to:

- Add or delete secure shell (SSH) public User keys to the `authorized_keys` file of user accounts that have been configured on the system, including the admin account. This allows authentication to user accounts using SSH keys rather than password challenge.
- Edit the following SSH server configuration settings:
  - Public Key Authentication Algorithms
  - Cipher Algorithms
  - KEX Algorithms
  - MAC Methods
  - Minimum Server Key Size.

**Note**

To configure Host keys, which are used when performing SCP pushes of log files from the Cisco appliance to other host machines, use `logconfig -> hostkeyconfig`. For more information, see [Chapter 38, “Logging.”](#)

**Note**

After using the `sshconfig` command, a reboot is required for changes to take effect.

Using `hostkeyconfig`, you can scan for keys of remote hosts and add them to the Cisco appliance.

**Related Topics**

- [Example: Install a New Public Key, page 32-28](#)
- [Example: Edit SSH Server Configuration, page 32-29](#)

## Example: Install a New Public Key

In the following example, a new public key is installed for the administrator account:

```
mail.example.com> sshconfig
```

Choose the operation you want to perform:

- SSHD - Edit SSH server settings.
  - USERKEY - Edit SSH User Key settings
- ```
[> userkey
```

Currently installed keys for admin:

Choose the operation you want to perform:

- NEW - Add a new key.
 - USER - Switch to a different user to edit.
- ```
[> new
```

Please enter the public SSH key for authorization.  
Press enter on a blank line to finish.  
[-paste public key for user authentication here-]

Choose the operation you want to perform:

```
- SSHD - Edit SSH server settings.
- USERKEY - Edit SSH User Key settings
[]>
```

## Example: Edit SSH Server Configuration

The following example shows how to edit the SSH server configuration.

```
mail.example.com> sshconfig
```

Choose the operation you want to perform:

```
- SSHD - Edit SSH server settings.
- USERKEY - Edit SSH User Key settings
[]> sshd
```

ssh server config settings:

Public Key Authentication Algorithms:

```
rsa1
ssh-dss
ssh-rsa
```

Cipher Algorithms:

```
aes128-ctr
aes192-ctr
aes256-ctr
arcfour256
arcfour128
aes128-cbc
3des-cbc
blowfish-cbc
cast128-cbc
aes192-cbc
aes256-cbc
arcfour
rijndael-cbc@lysator.liu.se
```

MAC Methods:

```
hmac-md5
hmac-sha1
umac-64@openssh.com
hmac-ripemd160
hmac-ripemd160@openssh.com
hmac-sha1-96
hmac-md5-96
```

Minimum Server Key Size:

```
1024
```

KEX Algorithms:

```
diffie-hellman-group-exchange-sha256
diffie-hellman-group-exchange-sha1
diffie-hellman-group14-sha1
diffie-hellman-group1-sha1
```

Choose the operation you want to perform:

```
- SETUP - Setup SSH server configuration settings
[]> setup
```

Enter the Public Key Authentication Algorithms do you want to use

```
[rsa1,ssh-dss,ssh-rsa]> rsa1
```

Enter the Cipher Algorithms do you want to use

```
[aes128-ctr,aes192-ctr,aes256-ctr,arcfour256,arcfour128,aes128-cbc,3des-cbc,blowfish-cbc,cast128-cbc,aes192-cbc,aes256-cbc,arcfour,rijndael-cbc@lysator.liu.se]> aes192-ctr
```

```

Enter the MAC Methods do you want to use
[hmac-md5,hmac-sha1,umac-64@openssh.com,hmac-ripemd160,hmac-ripemd160@openssh.com,hmac-sha1-96,hmac-md5-96]> hmac-sha1

Enter the Minimum Server Key Size do you want to use
[1024]> 2048

Enter the KEX Algorithms do you want to use
[diffie-hellman-group-exchange-sha256,diffie-hellman-group-exchange-sha1,diffie-hellman-group14-sha1,diffie-hellman-group1-sha1]> diffie-hellman-group-exchange-sha1

ssh server config settings:
Public Key Authentication Algorithms:
 rsa1
Cipher Algorithms:
 aes192-ctr
MAC Methods:
 hmac-sha1
Minimum Server Key Size:
 2048
KEX Algorithms:
 diffie-hellman-group-exchange-sha1

Choose the operation you want to perform:
- SETUP - Setup SSH server configuration settings
[]>

```

## Remote SSH Command Execution

The CLI allows commands to be run via remote SSH command execution. See [Appendix A, “AsyncOS Quick Reference Guide”](#) for a list of commands. For example, the following command can be run from a remote host unchallenged if an SSH public key has been configured for the admin account on the Cisco appliance:

```

ssh admin@mail3.example.com status

Enter "status detail" for more information.

Status as of: Mon Jan 20 17:24:15 2003

Last counter reset: Mon Jan 20 17:08:21 2003

System status: online

[rest of command deleted]

```

## Viewing Active Administrator Sessions

To view all administrators currently logged into the appliance, and information about their sessions, click **Options > Active Sessions** at the top right of the page.



## System Administration



### Note

Several of the features or commands described in this section will affect, or be affected by routing precedence. Please see [IP Addresses, Interfaces, and Routing, page B-3](#) for more information.

- [Management of the Appliance, page 33-1](#)
- [Feature Keys, page 33-5](#)
- [Managing the Configuration File, page 33-7](#)
- [Managing Disk Space, page 33-16](#)
- [Service Updates, page 33-17](#)
- [Setting Up to Obtain Upgrades and Updates, page 33-18](#)
- [Upgrading AsyncOS, page 33-25](#)
- [Enabling Remote Power Management, page 33-29](#)
- [Reverting to a Previous Version of AsyncOS, page 33-30](#)
- [Configuring the Return Address for Appliance Generated Messages, page 33-33](#)
- [Alerts, page 33-34](#)
- [Changing Network Settings, page 33-53](#)
- [Configuring SSL Settings, page 33-58](#)
- [System Time, page 33-59](#)
- [Customizing Your View, page 33-60](#)
- [Overriding Internet Explorer Compatibility Mode, page 33-62](#)

## Management of the Appliance

The following tasks allow you to easily manage the common functions within the appliance. The following operations and commands are described:

- shutdown
- reboot
- suspend
- offline

- `resume`
- `resetconfig`
- `version`
- `updateconfig`
- `upgrade`

## Shutting Down or Rebooting the Appliance

After you shut down or reboot, you may restart the appliance later without losing any messages in the delivery queue.

You can use the `shutdown` or `reboot` command in the CLI, or use the GUI:

### Procedure

- 
- Step 1** Select **System Administration > Shutdown/Suspend**.
  - Step 2** In the **System Operations** section, choose **Shutdown** or **Reboot** from the **Operation** drop-down list.
  - Step 3** Enter a number of seconds to wait to allow open connections to complete before forcing them to close. The default delay is thirty (30) seconds.
  - Step 4** Click **Commit**.
- 

## Suspending Email Receiving and Delivery

AsyncOS allows you to suspend receiving and delivering of emails. You can suspend:

- Receiving of emails on a particular listener or multiple listeners.
- Delivery of all emails or emails to a particular domain or multiple domains.

Use the `suspend` command in the CLI, or use the GUI:

### Procedure

- 
- Step 1** Select **System Administration > Shutdown/Suspend**.
  - Step 2** Suspend receiving of emails on a particular listener or multiple listeners.  
In the **Mail Operations** section, select the functions and/or listeners to suspend. If the appliance has multiple listeners, you can suspend email receiving on individual listeners.
  - Step 3** Suspend the delivery of all emails or emails to a particular domain or multiple domains. Depending on your requirements, do one of the following:
    - To suspend the delivery of all emails, in **Specify Domain(s)/Subdomain(s)** field, enter **ALL**, and press **Enter**.
    - To suspend the delivery of emails to a specific domain or subdomain, in **Specify Domain(s)/Subdomain(s)** field, enter the domain or subdomain name or IP address, and press **Enter**. Use comma-separated text to add multiple entries.

- Step 4** Enter number of seconds to wait to allow open connections to complete before forcing them to close. If there are no open connections, the system goes offline immediately. The default delay is 30 seconds.
- Step 5** Click **Commit**.
- 

**What To Do Next**

When you are ready to resume suspended services, see [Resuming Suspended Email Receiving and Delivery, page 33-3](#).

## Resuming Suspended Email Receiving and Delivery

Use the Shutdown/Suspend page or the `resume` command to resume the suspended receiving and delivery of emails.

**Procedure**

- 
- Step 1** Select **System Administration > Shutdown/Suspend**.
- Step 2** In the **Mail Operations** section, select the functions and/or listeners to resume. If the appliance has multiple listeners, you can resume email receiving on individual listeners.
- Step 3** Resume the delivery of all emails or emails to a particular domain or multiple domains. In **Specify Domain(s)/Subdomain(s)** field, click the close icon on the intended entry.
- Step 4** Click **Commit**.
- 

## Taking an Appliance Offline Using the CLI

If Cisco support tells you to do so, place Cisco IronPort AsyncOS into the offline state.

When the system is offline:

- Inbound email connections are not accepted.
- Outbound email delivery is halted.
- Log transfers are halted.
- The CLI remains accessible.

**Procedure**

- 
- Step 1** Use the `offline` command.
- Step 2** Specify the number of seconds to wait before forcing open connections to close.
-



## Resetting to Factory Defaults

When physically transferring the appliance, you may want to start with factory defaults. The Reset Configuration section of the System Administration > Configuration File page, or the `resetconfig` command, resets *all* AsyncOS configuration values to factory defaults. This command is extremely destructive, and it should only be used when you are transferring the unit or as a last resort to solving configuration issues. It is recommended you run the System Setup wizard or the `systemsetup` command after resetting the configuration.

**Note**

The `resetconfig` command only works when the appliance is in the offline state. When the `resetconfig` command completes, the appliance returns to the online state, even before you run the `systemsetup` command again. However, mail delivery will not be resumed; you will have to turn mail delivery back on.

**Caution**

The `resetconfig` command will return all network settings to factory defaults, potentially disconnecting you from the CLI, disabling services that you used to connect to the appliance (FTP, Telnet, SSH, HTTP, HTTPS), and even removing additional user accounts you created with the `userconfig` command. Do not use this command if you are not able to reconnect to the CLI using the Serial interface or the default settings on the Management port through the default Admin user account.

## The resetconfig Command

```
mail3.example.com> offline

Delay (seconds, minimum 30):

[30]> 45

Waiting for listeners to exit...

Receiving suspended.

Waiting for outgoing deliveries to finish...

Mail delivery suspended.

mail3.example.com> resetconfig

Are you sure you want to reset all configuration values? [N]> Y

All settings have been restored to the factory default.
```

## Displaying the Version Information for AsyncOS

To determine which version of AsyncOS is currently installed on your appliance, use the System Overview page from the Monitor menu in the GUI (see [System Status, page 28-30](#)), or use the `version` command in the CLI.

## Feature Keys

### Adding and Managing Feature Keys

For physical appliances, feature keys are specific to the serial number of the appliance and specific to the feature being enabled (you cannot re-use a key from one system on another system).

To work with feature keys in the CLI, use the `featurekey` command.

#### Procedure

- 
- Step 1** Select **System Administration > Feature Keys**.

**Step 2** Perform actions:

| To                                                                                   | Do This                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| View the status of active feature keys                                               | Look at the <b>Feature Keys for &lt;serial number&gt;</b> section.                                                                                                                                                                                      |
| View feature keys that have been issued for your appliance but are not yet activated | Look at the <b>Pending Activation</b> section.<br>If you have enabled automatic download and activation, feature keys will never appear in this list.                                                                                                   |
| Check for recently-issued feature keys                                               | Click the <b>Check for New Keys</b> button in the Pending Activation section.<br><br>This is useful if you have not enabled automatic download and activation of feature keys, or if you need to download feature keys before the next automatic check. |
| Activate an issued feature key                                                       | Select the key in the <b>Pending Activation</b> list and click <b>Activate Selected Keys</b> .                                                                                                                                                          |
| Add a new feature key                                                                | Use the <b>Feature Activation</b> section.                                                                                                                                                                                                              |

**Related Topics**

- [Automating Feature Key Download and Activation, page 33-6](#)
- [Cisco Email Security Virtual Appliance License, page 33-7](#)

## Automating Feature Key Download and Activation

You can set the appliance to automatically check for, download, and activate feature keys that are issued for this appliance.

**Procedure**

- 
- Step 1** Select **System Administration > Feature Key Settings**.
- Step 2** Click **Edit Feature Key Settings**.
- Step 3** To see frequency of checks for new feature keys, click the (?) help button.
- Step 4** Specify settings.
- Step 5** Submit and commit your changes.
- 

**Related Topics**

- [Adding and Managing Feature Keys, page 33-5](#)

## Expired Feature Keys

If the feature key for the feature you are trying to access (via the GUI) has expired, please contact your Cisco representative or support organization.

## Cisco Email Security Virtual Appliance License

To set up and license an Email Security Virtual appliance, see the *Cisco Content Security Virtual Appliance Installation Guide*. This document is available from the location specified in [Documentation, page 1-7](#).

**Note**

You cannot open a Technical Support tunnel or run the System Setup Wizard before installing the virtual appliance license.

## Virtual Appliance License Expiration

After the virtual appliance license expires, the appliance will continue to deliver mail without security services for 180 days. Security service updates do not occur during this period.

Alerts will be sent 180 days, 150 days, 120 days, 90 days, 60 days, 30 days, 15 days, 5 days, 1 day and 0 seconds before the license expires, and at the same intervals before the grace period ends. These alerts will be of type “System” at severity level “Critical.” To ensure that you receive these alerts, see [Adding Alert Recipients, page 33-36](#).

These alerts are also logged in the system log.

Individual feature keys may expire earlier than the virtual appliance license. You will also receive alerts when these approach their expiration dates.

**Related Topics**

- [Reverting AsyncOS on Virtual Appliances May Impact the License, page 33-31](#)

## Managing the Configuration File

All configuration settings within the appliance 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” both the CLI and the GUI for making configuration changes.

- You can upload entire configuration file via FTP access, or you can paste portions of or an entire configuration file directly into the CLI.
- 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.)

## Managing Multiple Appliances with XML Configuration Files

- You can download an existing configuration file from one appliance, make changes to it, and upload it to a different appliance. This lets you manage an installation of multiple appliances more easily. Currently you may not load configuration files from C/X-Series appliances onto an M-Series appliance.
- You can divide an existing configuration file downloaded from one appliance into multiple subsections. You can modify those sections that are common among all appliances (in a multiple appliance environment) and load them onto other appliances as the subsections are updated.

For example, you could use an appliance in a test environment for testing the Global Unsubscribe command. When you feel that you have configured the Global Unsubscribe list appropriately, you could then load the Global Unsubscribe configuration section from the test appliance to all of your production appliances.

## Managing Configuration Files Using the GUI

To use the GUI to manage configuration files on your appliance, click the Configuration File link on the System Administration tab.

The Configuration File page contains the following sections:

- **Current Configuration** - used to save and export the current configuration file.
- **Load Configuration** - used to load a complete or partial configuration file.
- **End-User Safelist/Blocklist Database (Spam Quarantine)** - For information, see [Using Safelists and Blocklists to Control Email Delivery Based on Sender, page 31-7](#) and [Backing Up and Restoring the Safelist/Blocklist, page 31-13](#).
- **Reset Configuration** - used to reset the current configuration back to the factory defaults (you should save your configuration prior to resetting it).

## Saving and Exporting the Current Configuration File

Using the **Current Configuration** section of the **System Administration > Configuration File** page, you can save the current configuration file to your local machine, save it on the appliance (placed in the `configuration` directory in the FTP/SCP root), or email it to the address specified.

The following information is not saved with the configuration file:

- Certificates used for secure communications with services used by the URL filtering feature.
- CCO User IDs and Contract ID saved on the Contact Technical Support page.

You can mask the user's passwords by clicking the **Mask passwords in the Configuration Files** checkbox. Masking a password causes the original, encrypted password to be replaced with "\*\*\*\*\*" in the exported or saved file. Please note, however, that configuration files with masked passwords cannot be loaded back into AsyncOS.

You can encrypt the user's passwords by clicking the **Encrypt passwords in the Configuration Files** checkbox. The following are the critical security parameters in the configuration file that will be encrypted.

- Certificate private keys
- RADIUS passwords
- LDAP bind passwords
- Local users' password hashes
- SNMP password
- DK/DKIM signing keys
- Outgoing SMTP authentication passwords
- PostX encryption keys
- PostX encryption proxy password
- FTP Push log subscriptions' passwords
- IPMI LAN password
- Updater server URLs



**Note** For enhanced security, if encryption of sensitive data in the appliance is enabled in FIPS mode, **Plain passwords in the Configuration Files** option is not displayed on the web interface.

## Loading a Configuration File

Use the Load Configuration section of the System Administration > Configuration File page to load new configuration information into the appliance. You can load information in one of three methods:

- Placing information in the `configuration` directory and uploading it.
- Uploading the configuration file directly from your local machine.
- Pasting configuration information directly into the GUI.

Configuration files with masked passwords cannot be loaded.



**Note** In cluster mode, you can either choose to load the configuration for a cluster or an appliance. For instructions to load cluster configuration, see [Loading a Configuration in Clustered Appliances](#), page 39-22.

Regardless of the method, you must include the following tags at the top of your configuration:

```
<?xml version="1.0" encoding="ISO-8859-1"?>

<!DOCTYPE config SYSTEM "config.dtd">

<config>

 ... your configuration information in valid XML

</config>
```

The closing `</config>` tag should follow your configuration information. The values in XML syntax are parsed and validated against the DTD (document type definition) located in the `configuration` directory on your appliance. The DTD file is named `config.dtd`. If validation errors are reported at the command line when you use the `loadconfig` command, the changes are not loaded. You can download the DTD to validate configuration files outside of the appliance before uploading them.

In either method, you can import an entire configuration file (the information defined between the highest level tags: `<config></config>`), or a *complete* and *unique* sub-section of the configuration file, as long as it contains the declaration tags (above) and is contained within the `<config></config>` tags.

“Complete” means that the entire start and end tags for a given subsection as defined by the DTD are included. For example, uploading or pasting this:

```
<?xml version="1.0" encoding="ISO-8859-1"?>

<!DOCTYPE config SYSTEM "config.dtd">

<config>

 <autosupport_enabled>0</autosu

</config>
```

will cause validation errors, while uploading. This, however:

```
<?xml version="1.0" encoding="ISO-8859-1"?>

<!DOCTYPE config SYSTEM "config.dtd">

<config>

 <autosupport_enabled>0</autosupport_enabled>

</config>
```

will not.

“Unique” means that the subsection of the configuration file being uploaded or pasted is not ambiguous for the configuration. For example, a system can have only one hostname, so uploading this (including the declarations and `<config></config>` tags):

```
<hostname>mail4.example.com</hostname>
```

is allowed. However, a system can have multiple listeners defined, each with different Recipient Access Tables defined, so uploading only this:

```
<rat>

 <rat_entry>

 <rat_address>ALL</rat_address>

 <access>RELAY</access>
```

```
</rat_entry>

</rat>
```

is considered ambiguous and is not allowed, even though it is “complete” syntax.

**Caution**

When uploading or pasting a configuration file or subsections of a configuration file, you have the potential to erase uncommitted changes that may be pending.

**Caution**

If disk space allocations in the configuration file are smaller than the amount of data currently stored on the appliance, the oldest data will be deleted to meet the quota specified in the configuration file.

### Empty vs. Omitted Tags

Use caution when uploading or pasting sections of configuration files. If you do not include a tag, then its value in the configuration is not modified when you load a configuration file. However, if you include an empty tag, then its configuration setting is cleared.

For example, uploading this:

```
<listeners></listeners>
```

will remove all listeners from the system!

**Caution**

When uploading or pasting subsections of a configuration file, you have the potential to disconnect yourself from the GUI or CLI and to destroy large amounts of configuration data. Do not disable services with this command if you are not able to reconnect to the appliance using another protocol, the Serial interface, or the default settings on the Management port. Also, do not use this command if you are unsure of the exact configuration syntax as defined by the DTD. Always back up your configuration data prior to loading a new configuration file.

### Note About Loading Passwords for Log Subscriptions

If you attempt to load a configuration file that contains a log subscription that requires a password (for example, one that will use FTP push), the `loadconfig` command does not warn you about the missing password. The FTP push will fail and alerts will be generated until you configure the correct password using the `logconfig` command.

### Note About Character Set Encoding

The “encoding” attribute of the XML configuration file must be “ISO-8859-1” regardless of the character set you may be using to manipulate the file offline. Note that the encoding attribute is specified in the file whenever you issue the `showconfig`, `saveconfig`, or `mailconfig` commands:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
```

Currently, only configuration files with this encoding can be loaded.



## Resetting the Current Configuration

Resetting the current configuration causes your appliance to revert back to the original factory defaults. You should save your configuration prior to resetting it. Resetting the configuration via this button in the GUI is not supported in a clustering environment.

See [Resetting to Factory Defaults, page 33-4](#).

## CLI Commands for Configuration Files

The following commands allow you to manipulate the configuration files:

- `showconfig`
- `mailconfig`
- `saveconfig`
- `loadconfig`
- `resetconfig` (See [Resetting to Factory Defaults, page 33-4](#).)

## The `showconfig`, `mailconfig`, and `saveconfig` Commands

For the configuration commands `showconfig`, `mailconfig`, and `saveconfig`, you are prompted to choose whether to include passwords in the file that will be mailed or displayed. Choosing not to include passwords will leave any password field blank. You can choose not to include passwords if you are concerned about security breaches. However, configuration files without passwords will fail when loaded using the `loadconfig` command. See [Note About Loading Passwords for Log Subscriptions, page 33-11](#).



### Note

When saving, showing, or mailing your configuration file if you choose to include passwords (answer yes to “Do you want to include passwords?”) the passwords are encrypted. However, the private keys and certificates are included in unencrypted PEM format.

The `showconfig` command prints the current configuration to the screen.

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

```
Do you want to include passwords? Please be aware that a configuration without passwords
will fail when reloaded with loadconfig.
```

```
<?xml version="1.0" encoding="ISO-8859-1"?>
```

```
<!DOCTYPE config SYSTEM "config.dtd">
```

```
<!--
```

```
Product: IronPort model number Messaging Gateway Appliance(tm)
```

Model Number: *model number*

Version: *version of AsyncOS installed*

Serial Number: *serial number*

Current Time: *current time and date*

*[The remainder of the configuration file is printed to the screen.]*

Use the `mailconfig` command to email the current configuration to a user. A configuration file in XML format named `config.xml` will be attached to the message.

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

Please enter the email address to which you want to send

the configuration file.

```
[> administrator@example.com
```

Do you want to include passwords? Please be aware that a configuration without passwords will fail when reloaded with `loadconfig`. [N]> **y**

The configuration file has been sent to `administrator@example.com`.

The `saveconfig` command saves the configuration file with a unique filename to the configuration directory on the appliance.

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

Do you want to include passwords? Please be aware that a configuration without passwords will fail when reloaded with `loadconfig`. [N]> **y**

File written on machine "mail3.example.com" to the location  
"/configuration/C360-420E874BB4B3C41C5C71-1419B58528A0-20120105T214041.xml".  
Configuration saved.

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

## The loadconfig Command

Use the `loadconfig` to load new configuration information into the appliance. You can load information in one of two methods:

- Placing information in the `configuration` directory and uploading it.
- Pasting configuration information directly into the CLI.

See [Loading a Configuration File, page 33-9](#) for more information.

## Uploading Configuration Changes Using the CLI

### Procedure

- 
- |               |                                                                                                                                                                                                                                     |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Outside of the CLI, ensure that you are able to access the <code>configuration</code> directory of the appliance. See <a href="#">Appendix A, “FTP, SSH, SCP, and Telnet Access”</a> for more information.                          |
| <b>Step 2</b> | Place an entire configuration file or subsection of a configuration file in the <code>configuration</code> directory of the appliance, or edit an existing configuration that was created from the <code>saveconfig</code> command. |
| <b>Step 3</b> | Within the CLI, use the <code>loadconfig</code> command to load the configuration file you placed in the directory from Step 2, or paste the text (XML syntax) directly into the CLI.                                               |
- 

In this example, a file named `changed.config.xml` is uploaded and the changes are committed:

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

```
1. Paste via CLI
```

```
2. Load from file
```

```
[1]> 2
```

```
Enter the name of the file to import:
```

```
[1]> changed.config.xml
```

```
Values have been loaded.
```

```
Be sure to run "commit" to make these settings active.
```

In this example, a new configuration file is pasted directly at the command line. (Remember to type Control-D on a blank line to end the paste command.) Then, the system setup wizard is used to change the default hostname, IP address, and default gateway information. (For more information, see [Using the System Setup Wizard, page 3-13](#).) Finally, the changes are committed.

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

1. Paste via CLI
2. Load from file

```
[1]> 1
```

Paste the configuration file now. Press CTRL-D on a blank line when done.

*[The configuration file is pasted until the end tag </config>. Control-D is entered on a separate line.]*

Values have been loaded.

Be sure to run "commit" to make these settings active.

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

*[The system setup wizard is run.]*

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

Please enter some comments describing your changes:

```
[]> pasted new configuration file and changed default settings via
systemsetup
```

Do you want to save the current configuration for rollback? [Y]> n

Changes committed: Fri May 23 11:42:12 2014 GMT

# Managing Disk Space

## (Virtual Appliances Only) Increasing Available Disk Space

For virtual appliances running ESXi 5.5 and VMFS 5, you can allocate more than 2TB of disk space. For appliances running ESXi 5.1, the limit is 2 TB.

To add disk space to the virtual appliance instance:



### Note

ESX does not support disk space reduction. See the VMWare documentation for information.

### Before You Begin

Carefully determine the disk space increase needed.

**Step 1** Bring down the Email Security appliance instance.

**Step 2** Increase disk space using utilities or administrative tools provided by VMWare.

See information about changing the virtual disk configuration in the VMWare documentation. At time of release, this information for ESXi 5.5 was available here:

<http://pubs.vmware.com/vsphere-55/index.jsp?topic=%2Fcom.vmware.vsphere.hostclient.doc%2FGUID-81629CAB-72FA-42F0-9F86-F8FD0DE39E57.html>.

**Step 3** Go to **System Administration > Disk Management** and verify that your change has taken effect.

## Allocating Disk Space Quotas

You can optimize disk usage by allocating disk space on the appliance among the features that your deployment uses.

| To                                                                                                                                                                       | Do This                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>View disk space quotas and current usage for each service</li> <li>Reallocate disk space on your appliance at any time</li> </ul> | Go to <b>System Administration &gt; Disk Management</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Manage data volume                                                                                                                                                       | <ul style="list-style-type: none"> <li>For reporting and tracking services and the spam quarantine, the oldest data will be deleted automatically.</li> <li>For Policy, Virus and Outbreak quarantines, the default action configured in the quarantine will be taken. See <a href="#">Default Actions for Automatically Processed Quarantined Messages</a>, page 30-4.</li> <li>For the Miscellaneous quota, you must first manually delete data to reduce usage below the new quota you will set. See <a href="#">Managing Disk Space for the Miscellaneous Quota</a>, page 33-17.</li> </ul> |

## Managing Disk Space for the Miscellaneous Quota

The Miscellaneous quota includes System data and User data. You cannot delete System data. User data that you can manage includes the following types of files:

| To Manage                                                                     | Do this                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Log files                                                                     | Go to <b>System Administration &gt; Log Subscriptions</b> and: <ul style="list-style-type: none"> <li>Look to see which log directories consume the most disk space.</li> <li>Verify that you need all of the log subscriptions that are being generated.</li> <li>Verify that the log level is no more verbose than necessary.</li> <li>If feasible, reduce the rollover file size.</li> </ul> |
| Packet captures                                                               | Go to <b>Help and Support</b> (near the upper right side of your screen) > <b>Packet Capture</b> .                                                                                                                                                                                                                                                                                              |
| Configuration files<br>(These files are unlikely to consume much disk space.) | FTP to the <code>/data/pub</code> directory on the appliance.<br>To configure FTP access to the appliance, see <a href="#">Appendix A, “FTP, SSH, SCP, and Telnet Access.”</a>                                                                                                                                                                                                                  |
| Quota size                                                                    | Go to <b>System Administration &gt; Disk Management</b> .                                                                                                                                                                                                                                                                                                                                       |

## Ensuring That You Receive Alerts About Disk Space

You will begin to receive system alerts at warning level when Miscellaneous disk usage reaches 75% of the quota. You should take action when you receive these alerts.

To ensure that you receive these alerts, see [Adding Alert Recipients, page 33-36](#).

## Disk Space and Centralized Management

Disk space management is available only in machine mode, not in group or cluster mode.

## Service Updates

The following services require updates for maximum effectiveness:

- Feature Keys
- McAfee Anti-Virus definitions
- PXE Engine
- Sophos Anti-Virus definitions
- IronPort Anti-Spam rules
- Outbreak Filters rules
- Time zone rules

- URL categories (Used for URL filtering features. For details, see [Future URL Category Set Changes, page 15-22](#))
- Enrollment client (Used for updating certificates needed for communication with cloud-based services used for URL filtering features. For information, see [About the Connection to Cisco Web Security Services, page 15-3](#).)

**Note**

Settings for the RSA Email DLP engine and content matching classifiers are handled on the **Security Services > RSA Email DLP** page. See [About Updating the DLP Engine and Content Matching Classifiers, page 17-39](#) for more information.

Service update settings are used for all services that receive updates except DLP updates. You cannot specify unique settings for any individual service except DLP updates.

To set up the network and the appliance to obtain these critical updates, see [Setting Up to Obtain Upgrades and Updates, page 33-18](#).

## Setting Up to Obtain Upgrades and Updates

- [Options for Distributing Upgrades and Updates, page 33-18](#)
- [Configuring Your Network to Download Upgrades and Updates from the Cisco Servers, page 33-19](#)
- [Configuring the Appliance for Upgrades and Updates in Strict Firewall Environments, page 33-19](#)
- [Upgrading and Updating from a Local Server, page 33-19](#)
- [Hardware and Software Requirements for Upgrading and Updating from a Local Server, page 33-20](#)
- [Hosting an Upgrade Image on a Local Server, page 33-21](#)
- [Configuring Server Settings for Downloading Upgrades and Updates, page 33-21](#)
- [Configuring Automatic Updates, page 33-23](#)
- [Configuring the Appliance to Verify the Validity of Updater Server Certificate, page 33-23](#)
- [Configuring the Appliance to Trust Proxy Server Communication, page 33-25](#)

## Options for Distributing Upgrades and Updates

There are several ways to distribute AsyncOS upgrade and update files to your appliances:


- Each appliance can download the files directly from the Cisco update servers. This is the default method.
- You can download the files from Cisco once, and then distribute them to your appliances from a server within your network. See [Upgrading and Updating from a Local Server, page 33-19](#).

To choose and configure a method, see [Configuring Server Settings for Downloading Upgrades and Updates, page 33-21](#).

## Configuring Your Network to Download Upgrades and Updates from the Cisco Servers

The appliance connects directly to the Cisco update servers to find and download upgrades and updates:

**Figure 33-1 Streaming Update Method**



Cisco update servers use dynamic IP addresses. If you have strict firewall policies, you may need to configure a static location instead. For more information, see [Configuring the Appliance for Upgrades and Updates in Strict Firewall Environments](#), page 33-19.

Create a firewall rule to allow downloading of upgrades from Cisco update servers on ports 80 and 443.

## Configuring the Appliance for Upgrades and Updates in Strict Firewall Environments

The Cisco IronPort upgrade and update servers use dynamic IP addresses. If you have strict firewall policies, you may need to configure a static location for updates and AsyncOS upgrades.

### Procedure

- 
- Step 1** Contact Cisco Customer support to obtain the static URL address.
  - Step 2** Create a firewall rule to allow downloading of upgrades and updates from the static IP address on port 80.
  - Step 3** Choose **Security Services > Service Updates**.
  - Step 4** Click **Edit Update Settings**.
  - Step 5** On the Edit Update Settings page, in the “Update Servers (images)” section, choose Local Update Servers and enter the static URL received in step 1 in the Base URL field for AsyncOS upgrades and McAfee Anti-Virus definitions.
  - Step 6** Verify that IronPort Update Servers is selected for the “Update Servers (list)” section.
  - Step 7** Submit and commit your changes.
- 

## Upgrading and Updating from a Local Server

You can download AsyncOS upgrade images to a local server and host upgrades from within your own network rather than obtaining upgrades directly from Cisco’s update servers. Using this feature, an upgrade image is downloaded via HTTP to any server in your network that has access to the Internet. If you choose to download the upgrade image, you can then configure an internal HTTP server (an “update manager”) to host the AsyncOS images to your appliances.




Use a local server if your appliance does not have access to the internet, or if your organization restricts access to mirror sites used for downloads. Downloading AsyncOS upgrades to each appliance from a local server is generally faster than downloading from the Cisco IronPort servers.

**Note**

Cisco recommends using a local server only for AsyncOS upgrades. If you use a local update server for security update images, the local server does not automatically receive security updates from Cisco IronPort, so the appliances in your network may not always have the most current security services.

**Figure 33-2 Remote Update Method**

**Procedure**

- Step 1** Configure a local server to retrieve and serve the upgrade files.
- Step 2** Download the upgrade files.
- Step 3** Configure the appliance to use the local server using either the Security Services > Service Updates page in the GUI or the `updateconfig` command in the CLI.
- Step 4** Upgrade the appliance using either the System Administration > System Upgrade page or the `upgrade` command in the CLI.

## Hardware and Software Requirements for Upgrading and Updating from a Local Server

For *downloading* AsyncOS upgrade and update files, you must have a system in your internal network that has:

- Internet access to the Cisco Systems update servers.
- A web browser (see [Browser Requirements, page 2-1](#)).

**Note**

For this release, if you need to configure a firewall setting to allow HTTP access to this address, you must configure it using the DNS name and not a specific IP address.

For *hosting* AsyncOS update files, you must have a server in your internal network that has:

- A web server — for example, Microsoft IIS (Internet Information Services) or the Apache open source server — which:
  - supports the display of directory or filenames in excess of 24 characters
  - has directory browsing enabled
  - is configured for anonymous (no authentication) or basic (“simple”) authentication
  - contains at least 350MB of free disk space for each AsyncOS update image

## Hosting an Upgrade Image on a Local Server

After setting up a local server, go to [http://updates.ironport.com/fetch\\_manifest.html](http://updates.ironport.com/fetch_manifest.html) to download a ZIP file of an upgrade image. To download the image, enter your serial number (for a physical appliance) or a VLN (for a virtual appliance) and the version number of the appliance. You will then be presented with a list of available upgrades. Click on the upgrade version that you want to download, and unzip the ZIP file in the root directory on the local server while keeping the directory structure intact. To use the upgrade image, configure the appliance to use the local server on the Edit Update Settings page (or use `updateconfig` in the CLI).

The local server also hosts an XML file that limits the available AsyncOS upgrades for the appliances on your network to the downloaded upgrade image. This file is called the “manifest.” The manifest is located in the `asyncos` directory of the upgrade image ZIP file. After unzipping the ZIP file in the root directory of the local server, enter the full URL for the XML file, including the filename, on the Edit Update Settings page (or use `updateconfig` in the CLI).

For more information about remote upgrades, please see the Knowledge Base or contact your Cisco Support provider.

## Updates Through a Proxy Server

The appliance is configured (by default) to connect directly to Cisco’s update servers to receive updates. This connection is made by HTTP on port 80 and the content is encrypted. If you do not want to open this port in your firewall, you can define a proxy server and specific port from which the appliance can receive updated rules.

If you choose to use a proxy server, you can specify an optional authentication and port.

**Note**

If you define a proxy server, it will *automatically* be used for all service updates that are configured to use a proxy server. There is no way to turn off the proxy server for updates to any individual service.

## Configuring Server Settings for Downloading Upgrades and Updates

Specify the server and connection information required to download upgrades and updates to your appliance.

You can use the same or different settings for AsyncOS upgrades and for service updates.

### Before You Begin

Determine whether the appliance will download upgrades and updates directly from Cisco, or whether you will host these images from a local server on your network instead. Then set up your network to support the method you choose. See all topics under [Setting Up to Obtain Upgrades and Updates](#), page 33-18.

### Procedure

- Step 1** Choose **Security Services > Service Updates**.
- Step 2** Click **Edit Update Settings**.
- Step 3** Enter options:

| Setting                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Update Servers (images)</b> | <p>Choose whether to download Cisco IronPort AsyncOS upgrade images and service updates from the Cisco IronPort update servers or from a local server on your network. The default is the Cisco IronPort update servers for both upgrades and updates.</p> <p>To use the same settings for upgrades and updates, enter information in the visible fields.</p> <p>If you choose a local update server, enter the base URL and port number for the servers used to download the upgrades and updates. If the server requires authentication, you can also enter a valid user name and password.</p> <p>To enter separate settings solely for AsyncOS upgrades and McAfee Anti-Virus definitions, click the <b>Click to use different settings for AsyncOS</b> link.</p> <p><b>Note</b> Cisco Intelligent Multi-Scan requires a second local server to download updates for third-party anti-spam rules.</p> |
| <b>Update Servers (lists)</b>  | <p>To ensure that only upgrades and updates that are appropriate to your deployment are available to each appliance, Cisco IronPort generates a manifest list of the relevant files.</p> <p>Choose whether to download the lists of available upgrades and service updates (the manifest XML files) from the Cisco IronPort update servers or from a local server on your network.</p> <p>There are separate sections for specifying servers for updates and for AsyncOS upgrades. The default for upgrades and updates is the Cisco IronPort update servers.</p> <p>If you choose local update servers, enter the full path to the manifest XML file for each list, including the file name and HTTP port number for the server. If you leave the port field blank, AsyncOS uses port 80. If the server requires authentication, enter a valid user name and password.</p>                               |

| Setting                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Automatic Updates</b>  | <p>Enable automatic updates and the update interval (how often the appliance checks for updates) for Sophos and McAfee Anti-Virus definitions, Cisco Anti-Spam rules, Cisco Intelligent Multi-Scan rules, PXE Engine updates, Outbreak Filter rules, and time zone rules.</p> <p>Include a trailing s, m, or h to indicate seconds, minutes, or hours. Enter 0 (zero) to disable automatic updates.</p> <p><b>Note</b> You can only turn on automatic updates for DLP using the <b>Security Services &gt; RSA Email DLP</b> page. However, you must enable automatic updates for all services first. See <a href="#">About Updating the DLP Engine and Content Matching Classifiers, page 17-39</a> for more information.</p> |
| <b>Interface</b>          | Choose which network interface to use when contacting the update servers for the listed security component updates. The available proxy data interfaces are shown. By default, the appliance selects an interface to use.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>HTTP Proxy Server</b>  | <p>An optional proxy server used for the services listed in the GUI.</p> <p>If you specify a proxy server, it will be used to update ALL services.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>HTTPS Proxy Server</b> | An optional proxy server using HTTPS. If you define the HTTPS proxy server, it will be used to update the services listed in the GUI.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

**Step 4** Submit and commit your changes.

## Configuring Automatic Updates

### Procedure

- Step 1** Navigate to the Security Services > Service Updates page, and click **Edit Update Settings**.
- Step 2** Select the check box to enable automatic updates.
- Step 3** Enter an update interval (time to wait between checks for updates). Add a trailing **m** for minutes and **h** for hours. The maximum update interval is 1 hour.

## Configuring the Appliance to Verify the Validity of Updater Server Certificate

Email security appliance can check the validity of Cisco updater server certificate every time the appliance communicates the updater server. If you configure this option and the verification fails, updates are not downloaded and the details are logged in Updater Logs.

Use the `updateconfig` command to configure this option. The following example shows how to configure this option.

```
mail.example.com> updateconfig
```

```
Service (images):
```

```
Update URL:
```

```

Feature Key updates http://downloads.ironport.com/asyncos
Timezone rules Cisco IronPort Servers
Enrollment Client Updates Cisco IronPort Servers
Support Request updates Cisco IronPort Servers
Cisco IronPort AsyncOS upgrades Cisco IronPort Servers

Service (list): Update URL:

Timezone rules Cisco IronPort Servers
Enrollment Client Updates Cisco IronPort Servers
Support Request updates Cisco IronPort Servers

Service (list): Update URL:

Cisco IronPort AsyncOS upgrades Cisco IronPort Servers

Update interval: 5m

Proxy server: not enabled

HTTPS Proxy server: not enabled

Choose the operation you want to perform:
- SETUP - Edit update configuration.
- VALIDATE_CERTIFICATES - Validate update server certificates
- TRUSTED_CERTIFICATES - Manage trusted certificates for updates
[> validate_certificates

Should server certificates from Cisco update servers be validated?
[Yes]>

Service (images): Update URL:

Feature Key updates http://downloads.ironport.com/asyncos
Timezone rules Cisco IronPort Servers
Enrollment Client Updates Cisco IronPort Servers
Support Request updates Cisco IronPort Servers
Cisco IronPort AsyncOS upgrades Cisco IronPort Servers

Service (list): Update URL:

Timezone rules Cisco IronPort Servers
Enrollment Client Updates Cisco IronPort Servers
Support Request updates Cisco IronPort Servers

Service (list): Update URL:

Cisco IronPort AsyncOS upgrades Cisco IronPort Servers

Update interval: 5m

Proxy server: not enabled

HTTPS Proxy server: not enabled

Choose the operation you want to perform:
- SETUP - Edit update configuration.
- VALIDATE_CERTIFICATES - Validate update server certificates
- TRUSTED_CERTIFICATES - Manage trusted certificates for updates

```

```
[]>
```

## Configuring the Appliance to Trust Proxy Server Communication

If you are using a non-transparent proxy server, you can add the CA certificate used to sign the proxy certificate to the appliance. By doing so, the appliance trusts the proxy server communication.

Use the `updateconfig` command to configure this option. The following example shows how to configure this option.

```
mail.example.com> updateconfig
...
...
...
Choose the operation you want to perform:
- SETUP - Edit update configuration.
- VALIDATE_CERTIFICATES - Validate update server certificates
- TRUSTED_CERTIFICATES - Manage trusted certificates for updates
[]> trusted_certificates

Choose the operation you want to perform:
- ADD - Upload a new trusted certificate for updates.
[]> add

Paste certificates to be trusted for secure updater connections, blank to quit
Trusted Certificate for Updater:
Paste cert in PEM format (end with '.'):
-----BEGIN CERTIFICATE-----
MMIICiDCCAFgGgAwIBAgIBATANBgkqhkiG9w0BAQUFADCBgDELMAkGA1UEBhmCSU4x
DDAKBgNVBAGTA0tBUjENM.....
-----END CERTIFICATE-----
.

Choose the operation you want to perform:
- ADD - Upload a new trusted certificate for updates.
- LIST - List trusted certificates for updates.
- DELETE - Delete a trusted certificate for updates.
[]>
```

## Upgrading AsyncOS

|               | Do This                                                                                                                                                                       | See                                                                   |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------|
| <b>Step 1</b> | If you have not yet done so, configure settings that apply to all update and upgrade downloads, and set up your network to support and optionally distribute these downloads. | <a href="#">Setting Up to Obtain Upgrades and Updates, page 33-18</a> |
| <b>Step 2</b> | Understand when an upgrade is available and determine whether to install it.                                                                                                  | <a href="#">Notifications of Available Upgrades, page 33-26</a>       |
| <b>Step 3</b> | Perform required and recommended tasks before each upgrade.                                                                                                                   | <a href="#">Preparing to Upgrade AsyncOS, page 33-26</a>              |
| <b>Step 4</b> | Perform the upgrade.                                                                                                                                                          | <a href="#">Downloading and Installing the Upgrade, page 33-27</a>    |

## About Upgrading Clustered Systems

If you are upgrading clustered machines, please see [Upgrading Machines in a Cluster](#), page 39-12.

## About Batch Commands for Upgrade Procedures

Batch commands for upgrade procedures are documented in the *Cisco AsyncOS CLI Reference Guide* at [http://www.cisco.com/en/US/products/ps10154/prod\\_command\\_reference\\_list.html](http://www.cisco.com/en/US/products/ps10154/prod_command_reference_list.html).

## Notifications of Available Upgrades

By default, users with administrator and technician privileges will see a notification at the top of the web interface when an AsyncOS upgrade is available for the appliance.

On clustered machines, actions apply only to the machine to which you are logged in.

| To                                                                                                                            | Do This                                                                                     |
|-------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| View more information about the latest upgrade                                                                                | Hover over the upgrade notification.                                                        |
| View a list of all available upgrades                                                                                         | Click the down arrow in the notification.                                                   |
| Dismiss a current notification.<br>The appliance will not display another notification until a new upgrade becomes available. | Click the down arrow, then select <b>Clear the notification</b> , then click <b>Close</b> . |
| Prevent future notifications (Users with Administrator privileges only.)                                                      | Go to <b>System Administration &gt; System Upgrade</b> .                                    |

## Notifications of Available Upgrades

By default, users with administrator and technician privileges will see a notification at the top of the web interface when an AsyncOS upgrade is available for the appliance.

On clustered machines, actions apply only to the machine to which you are logged in.

| To                                                                                                                            | Do This                                                                                     |
|-------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| View more information about the latest upgrade                                                                                | Hover over the upgrade notification.                                                        |
| View a list of all available upgrades                                                                                         | Click the down arrow in the notification.                                                   |
| Dismiss a current notification.<br>The appliance will not display another notification until a new upgrade becomes available. | Click the down arrow, then select <b>Clear the notification</b> , then click <b>Close</b> . |
| Prevent future notifications (Users with Administrator privileges only.)                                                      | Go to <b>System Administration &gt; System Upgrade</b> .                                    |

## Preparing to Upgrade AsyncOS

As a best practice, Cisco recommends preparing for an upgrade by taking the following steps.

### Procedure

- 
- Step 1** Save the XML configuration file off-box. If you need to revert to the pre-upgrade release for any reason, you will need this file.
- Step 2** If you are using the Safelist/Blocklist feature, export the list off-box.
- Step 3** Suspend all listeners. If you perform the upgrade from the CLI, use the `suspendlistener` command. If you perform the upgrade from the GUI, listener suspension occurs automatically.
- Step 4** Wait for the queue to empty. You can use the `workqueue` command to view the number of messages in the work queue or the `rate` command in the CLI to monitor the message throughput on your appliance.



**Note** Re-enable the listeners post-upgrade.

---

## Downloading and Installing the Upgrade

You can download and install in a single operation, or download in the background and install later.



### Note

When downloading and upgrading AsyncOS in a single operation from a local server instead of from a Cisco IronPort server, the upgrade installs immediately *while downloading*. A banner displays for 10 seconds at the beginning of the upgrade process. While this banner is displayed, you have the option to type Control-C to exit the upgrade process before downloading starts.

---

### Before You Begin

- Choose whether you will download upgrades directly from Cisco or will host upgrade images from a server on your network. Then set up your network to support the method you choose. Then configure the appliance to obtain upgrades from your chosen source. See [Setting Up to Obtain Upgrades and Updates, page 33-18](#) and [Configuring Server Settings for Downloading Upgrades and Updates, page 33-21](#).
- If you will install the upgrade now, follow the instructions in [Preparing to Upgrade AsyncOS, page 33-26](#).
- If you are installing the upgrade in a clustered system, see [Upgrading Machines in a Cluster, page 39-12](#).
- If you will only download the upgrade, there are no prerequisites until you are ready to install it.

### Procedure

- 
- Step 1** Choose **System Administration > System Upgrade**.
- Step 2** Click **Upgrade Options**.
- Step 3** Choose an option:



| To                                                     | Do This                                                                                                                                                                                                          |
|--------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Download and install the upgrade in a single operation | Click <b>Download and Install</b> .<br>If you have already downloaded an installer, you will be prompted to overwrite the existing download.                                                                     |
| Download an upgrade installer                          | Click <b>Download only</b> .<br>If you have already downloaded an installer, you will be prompted to overwrite the existing download.<br>The installer downloads in the background without interrupting service. |
| Install a downloaded upgrade installer                 | Click <b>Install</b> .<br>This option appears only if an installer has been downloaded.<br>The AsyncOS version to be installed is noted below the Install option.                                                |

**Step 4** Unless you are installing a previously-downloaded installer, select an AsyncOS version from the list of available upgrades.

**Step 5** If you are installing:

- a. Choose whether or not to save the current configuration to the `configuration` directory on the appliance.
- b. Choose whether or not to mask the passwords in the configuration file.



**Note** You cannot load a configuration file with masked passwords using the Configuration File page in the GUI or the `loadconfig` command in the CLI.

- c. If you want to email copies of the configuration file, enter the email addresses to which you want to email the file. Use commas to separate multiple email addresses.

**Step 6** Click **Proceed**.

**Step 7** If you are installing:

- a. Be prepared to respond to prompts during the process.  
The process pauses until you respond.  
A progress bar appears near the top of the page.
- b. At the prompt, click **Reboot Now**.
- c. After about 10 minutes, access the appliance again and log in.  
If you feel you need to power-cycle the appliance to troubleshoot an upgrade issue, do not do so until at least 20 minutes have passed since you rebooted.

#### What To Do Next

- If the process was interrupted, you must start the process again.
- If you downloaded but did not install the upgrade:

When you are ready to install the upgrade, follow these instructions from the beginning, including the prerequisites in the Before You Begin section, but choose the Install option.

- If you installed the upgrade:
  - Re-enable (resume) the listeners.
  - Save a configuration file for the new system. For information, see [Managing the Configuration File, page 33-7](#).
- After upgrade is complete, re-enable listeners.

## Viewing Status of, Canceling, or Deleting a Background Download

### Procedure

- Step 1** Choose **System Administration > System Upgrade**.
- Step 2** Click **Upgrade Options**.
- Step 3** Choose an option:

| To                            | Do This                                                                                                                                                                     |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| View download status          | Look in the middle of the page.<br><br>If there is no download in progress and no completed download waiting to be installed, you will not see download status information. |
| Cancel a download             | Click the <b>Cancel Download</b> button in the middle of the page.<br><br>This option appears only while a download is in progress.                                         |
| Delete a downloaded installer | Click the <b>Delete File</b> button in the middle of the page.<br><br>This option appears only if an installer has been downloaded.                                         |

- Step 4** (Optional) View the Upgrade Logs.

## Enabling Remote Power Management

The ability to remotely reset the power for the appliance chassis is available only on the following hardware: C380 and C680.

If you want to be able to remotely reset appliance power, you must enable and configure this functionality in advance, using the procedure described in this section.

### Before You Begin

- Cable the dedicated Remote Power Management port directly to a secure network. For information, see the Hardware Installation Guide.
- Ensure that the appliance is accessible remotely; for example, open any necessary ports through the firewall.

- This feature requires a unique IPv4 address for the dedicated Remote Power Management interface. This interface is configurable only via the procedure described in this section; it cannot be configured using the `ipconfig` command.
- In order to cycle appliance power, you will need a third-party tool that can manage devices that support the Intelligent Platform Management Interface (IPMI) version 2.0. Ensure that you are prepared to use such a tool.
- For more information about accessing the command-line interface, see the CLI reference guide.

#### Procedure

- 
- Step 1** Use SSH, telnet, or the serial console port to access the command-line interface.
- Step 2** Sign in using an account with Administrator access.
- Step 3** Enter the following commands:
- ```
remotepower
setup
```
- Step 4** Follow the prompts to specify the following:
- The dedicated IP address for this feature, plus netmask and gateway.
 - The username and password required to execute the power-cycle command.
- These credentials are independent of other credentials used to access your appliance.
- Step 5** Enter `commit` to save your changes.
- Step 6** Test your configuration to be sure that you can remotely manage appliance power.
- Step 7** Ensure that the credentials that you entered will be available to you in the indefinite future. For example, store this information in a safe place and ensure that administrators who may need to perform this task have access to the required credentials.
-

Related Topics

- [Remotely Resetting Appliance Power, page 40-27](#)

Reverting to a Previous Version of AsyncOS

AsyncOS includes the ability to revert the AsyncOS operating system to a previous qualified build for emergency uses.

Reversion Impact

Using the `revert` command on a appliance is a very destructive action. This command destroys all configuration logs and databases. Only the network information for the management interface is preserved--all other network configuration is deleted. In addition, reversion disrupts mail handling until the appliance is reconfigured. Because this command destroys network configuration, you may need physical local access to the appliance when you want to issue the `revert` command.

**Caution**

You must have a configuration file for the version you wish to revert to. Configuration files are *not* backwards-compatible.

Reverting AsyncOS on Virtual Appliances May Impact the License

If you revert from AsyncOS 9.0 for Email to AsyncOS 8.5 for Email, the license does not change.

If you revert from AsyncOS 9.0 for Email to AsyncOS 8.0 for Email, there is no longer a 180-day grace period during which the appliance delivers mail without security features.

Feature key expiration dates do not change in either case.

Related Topics

- [Virtual Appliance License Expiration, page 33-7](#)

Reverting AsyncOS

Procedure

Step 1 Ensure that you have the configuration file for the version you wish to revert to. Configuration files are not backwards-compatible. To do this, you can email the file to yourself or FTP the file. A simple way to do this is to run the `mailconfig` CLI command.

Step 2 Save a backup copy of the current configuration of your appliance (with passwords unmasked) on another machine.



Note This is not the configuration file you will load after reverting.

Step 3 If you use the Safelist/Blocklist feature, export the Safelist/Blocklist database to another machine.

Step 4 Wait for the mail queue to empty.

Step 5 Log into the CLI of the appliance you want to revert.

When you run the `revert` command, several warning prompts are issued. After these warning prompts are accepted, the revert action takes place immediately. Therefore, do not begin the reversion process until after you have completed the pre-reversion steps.

Step 6 From the CLI, Issue the `revert` command.



Note The reversion process is time-consuming. It may take fifteen to twenty minutes before reversion is complete and console access to the appliance is available again.

The following example shows the `revert` command:

```
mail.mydomain.com> revert
```

This command will revert the appliance to a previous version of AsyncOS.

WARNING: Reverting the appliance is extremely destructive.

The following data will be destroyed in the process:

- all configuration settings (including listeners)
- all log files
- all databases (including messages in Virus Outbreak and Policy quarantines)
- all reporting data (including saved scheduled reports)
- all message tracking data
- all IronPort Spam Quarantine message and end-user safelist/blocklist data

Only the network settings will be preserved.

Before running this command, be sure you have:

- saved the configuration file of this appliance (with passwords unmasked)
- exported the IronPort Spam Quarantine safelist/blocklist database to another machine (if applicable)
- waited for the mail queue to empty

Reverting the device causes an immediate reboot to take place.

After rebooting, the appliance reinitializes itself and reboots again to the desired version.

Do you want to continue?

Are you *really* sure you want to continue? yes

Available version	Install date
=====	=====

Available version	Install date
1. 5.5.0-236	Tue Aug 28 11:03:44 PDT 2007
2. 5.5.0-330	Tue Aug 28 13:06:05 PDT 2007
3. 5.5.0-418	Wed Sep 5 11:17:08 PDT 2007

Please select an AsyncOS version: 2

You have selected "5.5.0-330".

The system will now reboot to perform the revert operation.

- Step 7** Wait for the appliance to reboot twice.
- Step 8** After the machine reboots twice, use the serial console to configure an interface with an accessible IP address using the `interfaceconfig` command.
- Step 9** Enable FTP or HTTP on one of the configured interfaces.
- Step 10** Either FTP the XML configuration file you created, or paste it into the GUI interface.
- Step 11** Load the XML configuration file of the version you are reverting to.
- Step 12** If you use the Safelist/Blocklist feature, import and restore the Safelist/Blocklist database.
- Step 13** Commit your changes.
- The reverted appliance should now run using the selected AsyncOS version.

Configuring the Return Address for Appliance Generated Messages

You can configure the envelope sender for mail generated by AsyncOS for the following circumstances:

- Anti-Virus notifications
- Bounces
- DMARC feedback
- Notifications (`notify()` and `notify-copy()` filter actions)
- Quarantine notifications (and “Send Copy” in quarantine management)
- Reports
- All other messages

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.

You can modify the return address for system-generated email messages in the GUI or in the CLI using the `addressconfig` command.

Procedure

-
Step 1	Navigate to the System Administration > Return Addresses page.
Step 2	Click Edit Settings .
Step 3	Make changes to the address or addresses you want to modify
Step 4	Submit and commit your changes.
-

Alerts

Alert messages are automatically-generated standard email messages that contain information about events occurring on the 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 appliance. You can specify, at a much more granular level, which alert messages are sent to which users and for which severity of event they are sent. Manage alerts via the System Administration > Alerts page in the GUI (or via the `alertconfig` command in the CLI).

Alert Severities

Alerts can be sent for the following severities:

- Critical: Requires immediate attention.
- Warning: Problem or error requiring further monitoring and potentially immediate attention.
- Information: Information generated in the routine functioning of this device.

AutoSupport

To allow Cisco to better support and design future system changes, the appliance can be configured to send Cisco Systems a copy of all alert messages generated by the system. This feature, called AutoSupport, is a useful way to allow our team to be proactive in supporting your needs. AutoSupport also sends weekly reports noting the uptime of the system, the output of the `status` command, and the AsyncOS version used.

By default, alert recipients set to receive Information severity level alerts for System alert types will receive a copy of every message sent to Cisco. This can be disabled if you do not want to send the weekly alert messages internally. To enable or disable this feature, see [Configuring Alert Settings, page 33-37](#).

Alert Delivery

Alerts sent from the appliance to addresses specified in the Alert Recipient follow SMTP routes defined for those destinations.

Because alert messages can be used to inform you of problems within your appliance, they are not sent using AsyncOS's normal mail delivery system. Instead, alert messages pass through a separate and parallel email system designed to operate even in the face of significant system failure in AsyncOS.

The alert mail system does not share the same configuration as AsyncOS, which means that alert messages may behave slightly differently from other mail delivery:

- Alert messages are delivered using standard DNS MX and A record lookups.
 - They do not use `smtproutes` in AsyncOS versions older than 5.X.
 - They do cache the DNS entries for 30 minutes and the cache is refreshed every 30 minutes, so in case of DNS failure the alerts still go out.
- Alert messages do not pass through the work queue, so they are not scanned for viruses or spam. They are also not subjected to message filters or content filters.
- Alert messages do not pass through the delivery queue, so they are not affected by bounce profiles or destination control limits.

Example Alert Message

Date: 23 Mar 2005 21:10:19 +0000

To: joe@example.com

From: IronPort C60 Alert [alert@example.com]

Subject: Critical-example.com: (Anti-Virus) update via http://newproxy.example.com failed

The Critical message is:

update via http://newproxy.example.com failed

Version: 4.5.0-419

Serial Number: XXXXXXXXXXXX-XXXXXXX

Timestamp: Tue May 10 09:39:24 2005

For more information about this error, please see

<http://support.ironport.com>

If you desire further information, please contact your support provider.

Adding Alert Recipients

The alerting engine allows for granular control over which alerts are sent to which alert recipients. For example, you can configure the system to send only specific alerts to an alert recipient, configuring an alert recipient to receive notifications only when Critical (severity) information about the System (alert type) is sent.

**Note**

If you enabled AutoSupport during System Setup, the email address specified will receive alerts for all severities and classes by default. You can change this configuration at any time.

Procedure

- Step 1** Select **System Administration > Alerts**.
- Step 2** Click **Add Recipient**.
- Step 3** Enter the recipient's email address. You can enter multiple addresses, separated by commas.

- Step 4** (Optional) If you want to receive software release and critical support notification alerts from Cisco Support, check the **Release and Support Notifications** checkbox.
- Step 5** Select the alert types and severities that this recipient will receive.
- Step 6** Submit and commit your changes.
-

Configuring Alert Settings

The following settings apply to all alerts.



Note

Use the `alertconfig` CLI command to define the number of alerts to save on the appliance to view later.

Procedure

- Step 1** Click **Edit Settings** on the Alerts page.
- Step 2** Enter a Header From: address to use when sending alerts, or select Automatically Generated (“alert@<hostname>”).
- Step 3** Mark the checkbox if you want to specify the number of seconds to wait between sending duplicate alerts. For more information, see [Sending Duplicate Alerts, page 33-38](#).
- Specify the initial number of seconds to wait before sending a duplicate alert.
 - Specify the maximum number of seconds to wait before sending a duplicate alert.
- Step 4** You can enable AutoSupport by checking the IronPort AutoSupport option. For more information about AutoSupport, see [AutoSupport, page 33-34](#).
- If AutoSupport is enabled, the weekly AutoSupport report is sent to alert recipients set to receive System alerts at the Information level. You can disable this via the checkbox.
- Step 5** Submit and commit your changes.
-

Alert Settings

Alert settings control the general behavior and configuration of alerts, including:

- The RFC 2822 Header From: when sending alerts (enter an address or use the default “alert@<hostname>”). You can also set this via the CLI, using the `alertconfig -> from` command.
- The initial number of seconds to wait before sending a duplicate alert.
- The maximum number of seconds to wait before sending a duplicate alert.
- The status of AutoSupport (enabled or disabled).
- The sending of AutoSupport’s weekly status reports to alert recipients set to receive System alerts at the Information level.

Sending Duplicate Alerts

You can specify the initial number of seconds to wait before AsyncOS will send a duplicate alert. If you set this value to 0, duplicate alert summaries are not sent and instead, all duplicate alerts are sent without any delay (this can lead to a large amount of email over a short amount of time). The number of seconds to wait between sending duplicate alerts (alert interval) is increased after each alert is sent. The increase is the number of seconds to wait plus twice the last interval. So a 5 second wait would have alerts sent at 5 seconds, 15 seconds, 35 seconds, 75 seconds, 155 seconds, 315 seconds, etc.

Eventually, the interval could become quite large. You can set a cap on the number of seconds to wait between intervals via the maximum number of seconds to wait before sending a duplicate alert field. For example, if you set the initial value to 5 seconds, and the maximum value to 60 seconds, alerts would be sent at 5 seconds, 15 seconds, 35 seconds, 60 seconds, 120 seconds, etc.

Viewing Recent Alerts

The Email Security appliances saves the latest alerts so you can view them in both the GUI and the CLI in case you lose or delete the alert messages. These alerts cannot be downloaded from the appliance.

To view a list of the latest alerts, click the **View Top Alerts** button on the Alerts page or use the `displayalerts` command in the CLI. You can arrange the alerts in the GUI by date, level, class, text, and recipient.

By default, the appliance saves a maximum of 50 alerts to displays in the **Top Alerts** window. Use the `alertconfig -> setup` command in the CLI to edit the number of alerts that the appliance saves. If you want to disable this feature, change the number of alerts to 0.

Alert Descriptions

The following tables list alerts by classification, including the alert name (internal descriptor used by Cisco), actual text of the alert, description, severity (critical, information, or warning) and the parameters (if any) included in the text of the message. The value of the parameter is replaced in the actual text of the alert. For example, an alert message below may mention “\$ip” in the message text. “\$ip” is replaced by the actual IP address when the alert is generated.

Anti-Spam Alerts

[Table 33-1](#) contains a list of the various anti-spam alerts that can be generated by AsyncOS, including a description of the alert and the alert severity.

Table 33-1 Listing of Possible Anti-Spam Alerts

Alert Name	Message and Description	Parameters
AS.SERVER.ALERT	\$engine anti-spam - \$message \$tb	'engine' - The type of anti-spam engine. 'message' - The log message. 'tb' - Traceback of the event.
	Critical. Sent when the anti-spam engine fails.	

Table 33-1 Listing of Possible Anti-Spam Alerts (continued)

Alert Name	Message and Description	Parameters
AS.TOOL.INFO_ALERT	Update - \$engine - \$message	'engine' - The anti-spam engine name 'message' - The message
	Information. Sent when there is a problem with the anti-spam engine.	
AS.TOOL.ALERT	Update - \$engine - \$message	'engine' - The anti-spam engine name 'message' - The message
	Critical. Sent when an update is aborted due to a problem with one of the tools used to manage the anti-spam engine.	

Anti-Virus Alerts

Table 33-2 contains a list of the various Anti-Virus alerts that can be generated by AsyncOS, including a description of the alert and the alert severity.

Table 33-2 Listing of Possible Anti-Virus Alerts

Alert Name	Message and Description	Parameters
AV.SERVER.ALERT / AV.SERVER.CRITICAL	\$engine antivirus - \$message \$tb	'engine' - The type of anti-virus engine. 'message' - The log message. 'tb' - Traceback of the event.
	Critical. Sent when there is a critical problem with the anti-virus scanning engine.	
AV.SERVER.ALERT.INFO	\$engine antivirus - \$message \$tb	'engine' - The type of anti-virus engine. 'message' - The log message. 'tb' - Traceback of the event.
	Information. Sent when an informational event occurs with the anti-virus scanning engine.	
AV.SERVER.ALERT.WARN	\$engine antivirus - \$message \$tb	'engine' - The type of anti-virus engine. 'message' - The log message. 'tb' - Traceback of the event.
	Warning. Sent when there is a problem with the anti-virus scanning engine.	
MAIL.ANTIVIRUS. ERROR_MESSAGE	MID \$mid antivirus \$what error \$tag	'mid' - MID 'what' - The error that happened. 'tag' - Virus outbreak name if set.
	Critical. Sent when anti-virus scanning produces an error while scanning a message.	

Table 33-2 Listing of Possible Anti-Virus Alerts (continued)

Alert Name	Message and Description	Parameters
MAIL.SCANNER. PROTOCOL_MAX_RETRY	MID \$mid is malformed and cannot be scanned by \$engine. Critical. The scanning engine attempted to scan the message unsuccessfully because the message is malformed. The maximum number of retries has been exceeded, and the message will be processed without being scanned by this engine.	'mid' - MID 'engine' - The engine being used

Directory Harvest Attack Prevention (DHAP) Alerts

Table 33-3 contains a list of the various DHAP alerts that can be generated by AsyncOS, including a description of the alert and the alert severity.

Table 33-3 Listing of Possible Directory Harvest Attack Prevention Alerts

Alert Name	Message and Description	Parameters
LDAP.DHAP_ALERT	LDAP: Potential Directory Harvest Attack detected. See the system mail logs for more information about this attack. Warning. Sent when a possible directory harvest attack is detected.	

Hardware Alerts

Table 33-4 contains a list of the various Hardware alerts that can be generated by AsyncOS, including a description of the alert and the alert severity.

Table 33-4 Listing of Possible Hardware Alerts

Alert Name	Message and Description	Parameters
INTERFACE.ERRORS	Port \$port: has detected \$in_err input errors, \$out_err output errors, \$col collisions please check your media settings. Warning. Sent when interface errors are detected.	'port' - Interface name. 'in_err' - The number of input errors since the last message. 'out_err' - The number of output errors since the last message. 'col' - The number of packet collisions since the last message.
MAIL.MEASUREMENTS_ FILESYSTEM	The \$file_system partition is at \$capacity% capacity Warning. Sent when a disk partition is nearing capacity (75%).	'file_system' - The name of the filesystem 'capacity' - How full the filesystem is in percent.
MAIL.MEASUREMENTS_ FILESYSTEM.CRITICAL	The \$file_system partition is at \$capacity% capacity Critical. Sent when a disk partition reaches 90% capacity (and at 95%, 96%, 97%, etc.).	'file_system' - The name of the filesystem 'capacity' - How full the filesystem is in percent.

Table 33-4 Listing of Possible Hardware Alerts (continued)

Alert Name	Message and Description	Parameters
SYSTEM.RAID_EVENT_ALERT	A RAID-event has occurred: \$error	'error' - The text of the RAID error.
	Warning. Sent when a critical RAID-event occurs.	
SYSTEM.RAID_EVENT_ALERT_INFO	A RAID-event has occurred: \$error	'error' - The text of the RAID error.
	Information. Sent when a RAID-event occurs.	

Spam Quarantine Alerts

Table 33-5 contains a list of the various spam quarantine alerts that can be generated by AsyncOS, including a description of the alert and the alert severity.

Table 33-5 Listing of Possible Spam Quarantine Alerts

Alert Name	Message and Description	Parameters
ISQ.CANNOT_CONNECT_OFF_BOX	ISQ: Could not connect to off-box quarantine at \$host:\$port	'host' - address of off-box quarantine 'port' - port to connect to on off-box quarantine
	Information. Sent when AsyncOS was unable to connect to the (off-box) IP address.	
ISQ.CRITICAL	ISQ: \$msg	'msg' - message to be displayed
	Critical. Sent when a critical spam quarantine error is encountered.	
ISQ.DB_APPROACHING_FULL	ISQ: Database over \$threshold% full	'threshold' - the percent full threshold at which alerting begins
	Warning. Sent when the spam quarantine database is nearly full.	
ISQ.DB_FULL	ISQ: database is full	
	Critical. Sent when the spam quarantine database is full.	
ISQ.MSG_DEL_FAILED	ISQ: Failed to delete MID \$mid for \$rcpt: \$reason	'mid' - MID 'rcpt' - Recipient or "all" 'reason' - Why the message was not deleted
	Warning. Sent when an email is not successfully deleted from the spam quarantine.	
ISQ.MSG_NOTIFICATION_FAILED	ISQ: Failed to send notification message: \$reason	'reason' - Why the notification was not sent
	Warning. Sent when a notification message is not successfully sent.	
ISQ.MSG_QUAR_FAILED		
	Warning. Sent when a message is not successfully quarantined.	
ISQ.MSG_RLS_FAILED	ISQ: Failed to release MID \$mid to \$rcpt: \$reason	'mid' - MID 'rcpt' - Recipient or "all" 'reason' - Why the message was not released
	Warning. Sent when a message is not successfully released.	

Table 33-5 *Listing of Possible Spam Quarantine Alerts (continued)*

Alert Name	Message and Description	Parameters
ISQ.MSG_RLS_FAILED_UNK_RCPTS	ISQ: Failed to release MID \$mid: \$reason	'mid' - MID
	Warning. Sent when a message is not successfully released because the recipient is unknown.	'reason' - Why the message was not released
ISQ.NO_EU_PROPS	ISQ: Could not retrieve \$user's properties. Setting defaults	'user' - end user name
	Information. Sent when AsyncOS is unable to retrieve information about a user.	
ISQ.NO_OFF_BOX_HOST_SET	ISQ: Setting up off-box ISQ without setting host	
	Information. Sent when AsyncOS is configured to reference an external quarantine, but the external quarantine is not defined.	

Safelist/Blocklist Alerts

contains a list of the various Safelist/Blocklist alerts that can be generated by AsyncOS, including a description of the alert and the alert severity

Table 33-6 *Listing of Possible Safelist/Blocklist Alerts*

Alert Name	Message and Description	Parameters
SLBL.DB.RECOVERY_FAILED	SLBL: Failed to recover End-User Safelist/Blocklist database: '\$error'.	'error' - error reason
	Critical. Failed to recover the Safelist/Blocklist database.	
SLBL.DB.SPACE_LIMIT	SLBL: End-User Safelist/Blocklist database exceeded allowed disk space: \$current of \$limit.	'current' - how much it has used, in MB
	Critical. The safelist/blocklist database exceeded the allowed disk space.	'limit' - the configured limit, in MB

System Alerts

[Table 33-7](#) contains a list of the various System alerts that can be generated by AsyncOS, including a description of the alert and the alert severity.

Table 33-7 *Listing of Possible System Alerts*

Component/Alert Name	Message and Description	Parameters
AMP.ENGINE.ALERT	See Ensuring That You Receive Alerts, page 16-11 in Chapter 16, "Ensuring That You Receive Alerts."	-
AsyncOS API Alerts	See "Alerts" section in the Cisco AsyncOS API for Email - Getting Started Guide .	-
COMMON.APP_FAILURE	An application fault occurred: \$error	'error' - The text of the error, typically a traceback.
	Warning. Sent when there is an unknown application failure.	

Table 33-7 Listing of Possible System Alerts (continued)

Component/Alert Name	Message and Description	Parameters
COMMON.KEY_EXPIRED_ALERT	Your "\$feature" key has expired. Please contact your authorized Cisco sales representative.	'feature' - The name of the feature that is about to expire.
	Warning. Sent when a feature key has expired.	
COMMON.KEY_EXPIRING_ALERT	Your "\$feature" key will expire in under \$days day(s). Please contact your authorized Cisco sales representative.	'feature' - The name of the feature that is about to expire. 'days' - The number of days it will expire.
	Warning. Sent when a feature key is about to expire.	
COMMON.KEY_FINAL_EXPIRING_ALERT	This is a final notice. Your "\$feature" key will expire in under \$days day(s). Please contact your authorized Cisco sales representative.	'feature' - The name of the feature that is about to expire. 'days' - The number of days it will expire.
	Warning. Sent as a final notice that a feature key is about to expire.	
KEYS.GRACE_EXPIRING_ALERT	All security services licenses for this Cisco Email Security Appliance have expired. The appliance will continue to deliver mail without security services for \$days days.	'days' - The number of days remaining in the grace period at the time the alert was sent. For more information about the grace period, see Virtual Appliance License Expiration , page 33-7.
	To renew security services licenses, Please contact your authorized Cisco sales representative.	
	Critical. Sent periodically from the start of the grace period for virtual appliance license expiration.	
KEYS.GRACE_FINAL_EXPIRING_ALERT	This is the final notice. All security services licenses for this Cisco Email Security Appliance have expired. The appliance will continue to deliver mail without security services for 1 day.	For more information about the grace period, see Virtual Appliance License Expiration , page 33-7.
	To renew security services licenses, Please contact your authorized Cisco sales representative.	
	Critical. Sent one day before the virtual appliance license expires.	
KEYS.GRACE_EXPIRED_ALERT	Your grace period has expired. All security service have expired, and your appliance is non-functional. The appliance will no longer deliver mail until a new license is applied.	For more information about the grace period, see Virtual Appliance License Expiration , page 33-7.
	To renew security services licenses, Please contact your authorized Cisco sales representative.	
	Critical. Sent when the grace period for virtual appliances has expired.	
DNS.BOOTSTRAP_FAILED	Failed to bootstrap the DNS resolver. Unable to contact root servers.	
	Warning. Sent when the appliance is unable to contact the root DNS servers.	
INTERFACE.FAILOVER.FAILURE_BACKUP_DETECTED	Standby port \$port on \$pair_name failure	'port' - Detected port 'pair_name' - Failover pair name.
	Warning. Sent when a backup NIC pairing interface fails.	

Table 33-7 Listing of Possible System Alerts (continued)

Component/Alert Name	Message and Description	Parameters
INTERFACE.FAILOVER.FAILURE_BACKUP_RECOVERED	Standby port \$port on \$pair_name okay	'port' - Failed port
	Information. Sent when a NIC pair failover is recovered.	'pair_name' - Failover pair name.
INTERFACE.FAILOVER.FAILURE_DETECTED	Port \$port failure on \$pair_name, switching to \$port_other	'port' - Failed port.
	Critical. Sent when a NIC pairing failover is detected due to an interface failure.	'port_other' - New port. 'pair_name' - Failover pair name.
INTERFACE.FAILOVER.FAILURE_DETECTED_NO_BACKUP	Port \$port_other on \$pair_name is down, can't switch to \$port_other	'port' - Failed port.
	Critical. Sent when a NIC pairing failover is detected due to an interface failure, but a backup interface is not available.	'port_other' - New port. 'pair_name' - Failover pair name.
INTERFACE.FAILOVER.FAILURE_RECOVERED	Recovered network on \$pair_name using port \$port	'port' - Failed port
	Information. Sent when a NIC pair failover is recovered.	'pair_name' - Failover pair name.
INTERFACE.FAILOVER.MANUAL	Manual failover to port \$port on \$pair_name	'port' - New active port.
	Information. Sent when a manual failover to another NIC pair is detected.	'pair_name' - Failover pair name.
COMMON.INVALID_FILTER	Invalid \$class: \$error	'class' - Either "Filter", "SimpleFilter", etc.
	Warning. Sent when an invalid filter is encountered.	'error' - Additional why-filter-is-invalid info.
IPBLOCKD.HOST_ADDED_TO_WHITELIST	The host at \$ip has been added to the blacklist because of an SSH DOS attack.	'ip' - IP address from which a login attempt occurred.
IPBLOCKD.HOST_ADDED_TO_BLACKLIST	The host at \$ip has been permanently added to the ssh whitelist.	
IPBLOCKD.HOST_REMOVED_FROM_BLACKLIST	The host at \$ip has been removed from the blacklist	
	Warning. IP addresses that try to connect to the appliance over SSH but do not provide valid credentials are added to the SSH blacklist if more than 10 failed attempts occur within two minutes. When a user logs in successfully from the same IP address, that IP address is added to the whitelist. Addresses on the whitelist are allowed access even if they are also on the blacklist. Entries are automatically removed from the blacklist after about a day.	
LDAP.GROUP_QUERY_FAILED_ALERT	LDAP: Failed group query \$name, comparison in filter will evaluate as false	'name' - The name of the query.
	Critical. Sent when an LDAP group query fails.	

Table 33-7 Listing of Possible System Alerts (continued)

Component/Alert Name	Message and Description	Parameters
LDAP.HARD_ERROR	LDAP: work queue processing error in \$name reason \$why	'name' - The name of the query. 'why' - Why the error happened.
	Critical. Sent when an LDAP query fails completely (after trying all servers).	
LOG.ERROR.*	Critical. Various logging errors.	
MAIL.FILTER.RULE_MATCH_ALERT	MID \$mid matched the \$rule_name rule. \n Details: \$details	'mid' - Unique identification number of the message. 'rule_name' - The name of the rule that matched. 'details' - More information about the message or the rule.
	Information. Sent every time when a Header Repeats rule evaluates to true.	
MAIL.PERRCPT.LDAP_GROUP_QUERY_FAILED	LDAP group query failure during per-recipient scanning, possible LDAP misconfiguration or unreachable server.	
	Critical. Sent when an LDAP group query fails during per-recipient scanning.	
MAIL.QUEUE.ERROR.*	Critical. Various mail queue hard errors.	
MAIL.RES_CON_START_ALERT.MEMORY	This system (hostname: \$hostname) has entered a 'resource conservation' mode in order to prevent the rapid depletion of critical system resources. RAM utilization for this system has exceeded the resource conservation threshold of \$memory_threshold_start%. The allowed receiving rate for this system will be gradually decreased as RAM utilization approaches \$memory_threshold_halt%.	'hostname' - The name of the host. 'memory_threshold_start' - The percent threshold where memory tarpitting starts. 'memory_threshold_halt' - The percent threshold where the system will halt due to memory being too full.
	Critical. Sent when RAM utilization has exceeded the system resource conservation threshold.	
MAIL.RES_CON_START_ALERT.QUEUE_SLOW	This system (hostname: \$hostname) has entered a 'resource conservation' mode in order to prevent the rapid depletion of critical system resources. The queue is overloaded and is unable to maintain the current throughput.	'hostname' - The name of the host.
	Critical. Sent when the mail queue is overloaded and system resource conservation is enabled.	

Table 33-7 Listing of Possible System Alerts (continued)

Component/Alert Name	Message and Description	Parameters
MAIL.RES_CON_START_ALERT.QUEUE	This system (hostname: \$hostname) has entered a ‘resource conservation’ mode in order to prevent the rapid depletion of critical system resources. Queue utilization for this system has exceeded the resource conservation threshold of \$queue_threshold_start%. The allowed receiving rate for this system will be gradually decreased as queue utilization approaches \$queue_threshold_halt%.	‘hostname’ - The name of the host. ‘queue_threshold_start’ - The percent threshold where queue tarpitting starts. ‘queue_threshold_halt’ - The percent threshold where the system will halt due to the queue being too full.
	Critical. Sent when queue utilization has exceeded the system resource conservation threshold.	
MAIL.RES_CON_START_ALERT.WORKQ	This system (hostname: \$hostname) has entered a ‘resource conservation’ mode in order to prevent the rapid depletion of critical system resources. Listeners have been suspended because the current work queue size has exceeded the threshold of \$suspend_threshold. Listeners will be resumed once the work queue size has dropped to \$resume_threshold. These thresholds may be altered via use of the ‘tarpit’ command on the system CLI.	‘hostname’ - The name of the host. ‘suspend_threshold’ - Work queue size above which listeners are suspended. ‘resume_threshold’ - Work queue size below which listeners are resumed.
	Information. Sent when listeners are suspended because the work queue size is too big.	
MAIL.RES_CON_START_ALERT	This system (hostname: \$hostname) has entered a ‘resource conservation’ mode in order to prevent the rapid depletion of critical system resources.	‘hostname’ - The name of the host.
	Critical. Sent when the appliance enters “resource conservation” mode.	
MAIL.RES_CON_STOP_ALERT	This system (hostname: \$hostname) has exited ‘resource conservation’ mode as resource utilization has dropped below the conservation threshold.	‘hostname’ - The name of the host.
	Information. Sent when the appliance leaves ‘resource conservation’ mode.	
MAIL.SDS.CATEGORY_CHANGE	See Future URL Category Set Changes, page 15-22 .	—
MAIL.SDS.CERTIFICATE_INVALID	See Troubleshooting URL Filtering, page 15-10 .	
MAIL.SDS.ERROR_FETCHING_CERTIFICATE		
MAIL.WORK_QUEUE_PAUSED_NATURAL	work queue paused, \$num msgs, \$reason	‘num’ - The number of messages in the work queue. ‘reason’ - The reason the work queue is paused.
	Critical. Sent when the work queue is paused.	
MAIL.WORK_QUEUE_UNPAUSED_NATURAL	work queue resumed, \$num msgs	‘num’ - The number of messages in the work queue.
	Critical. Sent when the work queue is resumed.	
NTP.NOT_ROOT	Not running as root, unable to adjust system time	
	Warning. Sent when the appliance is unable to adjust time because NTP is not running as root.	

Table 33-7 Listing of Possible System Alerts (continued)

Component/Alert Name	Message and Description	Parameters
QUARANTINE.ADD_DB_ERROR	Unable to quarantine MID \$mid - quarantine system unavailable	'mid' - MID
	Critical. Sent when a message cannot be sent to a quarantine.	
QUARANTINE.DB_UPDATE_FAILED	Unable to update quarantine database (current version: \$version; target \$target_version)	'version' - The schema version detected.
	Critical. Sent when a quarantine database cannot be updated.	'target_version' - The target schema version.
QUARANTINE.DISK_SPACE_LOW	The quarantine system is unavailable due to a lack of space on the \$file_system partition.	'file_system' - The name of the filesystem.
	Critical. Sent when the disk space for quarantines is full.	
QUARANTINE.THRESHOLD_ALERT	Quarantine "\$quarantine" is \$full% full	'quarantine' - The name of the quarantine.
	Warning. Sent when a quarantine reaches 5%, 50%, or 75% of capacity.	'full' - The percentage of how full the quarantine is.
QUARANTINE.THRESHOLD_ALERT.SERIOUS	Quarantine "\$quarantine" is \$full% full	'quarantine' - The name of the quarantine.
	Critical. Sent when a quarantine reaches 95% of capacity.	
REPORTD.DATABASE_OPEN_FAILED_ALERT	The reporting system has encountered a critical error while opening the database. In order to prevent disruption of other services, reporting has been disabled on this machine. Please contact customer support to have reporting enabled. The error message is: \$err_msg	'err_msg' - The error message raised
	Critical. Sent if the reporting engine is unable to open the database.	
REPORTD.AGGREGATION_DISABLED_ALERT	Processing of collected reporting data has been disabled due to lack of logging disk space. Disk usage is above \$threshold percent. Recording of reporting events will soon become limited and reporting data may be lost if disk space is not freed up (by removing old logs, etc.). Once disk usage drops below \$threshold percent, full processing of reporting data will be restarted automatically.	'threshold' - The threshold value
	Warning. Sent if the system runs out of disk space. When the disk usage for a log entry exceeds the log usage threshold, reportd disables aggregation and sends the alert.	
REPORTING.CLIENT.UPDATE_FAILED_ALERT	Reporting Client: The reporting system has not responded for an extended period of time (\$duration).	'duration' - Length of time the client has been trying to contact the reporting daemon. This is a string in a human readable format ('1h 3m 27s').
	Warning. Sent if the reporting engine was unable to save reporting data.	

Table 33-7 Listing of Possible System Alerts (continued)

Component/Alert Name	Message and Description	Parameters
REPORTING.CLIENT. JOURNAL.FULL	Reporting Client: The reporting system is unable to maintain the rate of data being generated. Any new data generated will be lost.	
	Critical. Sent if the reporting engine is unable to store new data.	
REPORTING.CLIENT. JOURNAL.FREE	Reporting Client: The reporting system is now able to handle new data.	
	Information. Sent when the reporting engine is again able to store new data.	
PERIODIC_REPORTS. REPORT_TASK.BUILD_ FAILURE	A failure occurred while building periodic report '\$report_title'. This subscription has been removed from the scheduler.	'report_title' - the report title
	Critical. Sent when the reporting engine is unable to build a report.	
PERIODIC_REPORTS. REPORT_TASK.EMAIL_ FAILURE	A failure occurred while emailing periodic report '\$report_title'. This subscription has been removed from the scheduler.	'report_title' - the report title
	Critical. Sent when a report could not be emailed.	
PERIODIC_REPORTS. REPORT_TASK.ARCHIVE_FAILURE	A failure occurred while archiving periodic report '\$report_title'. This subscription has been removed from the scheduler.	'report_title' - the report title
	Critical. Sent when a report could not be archived.	
SENDERBASE.ERROR	Error processing response to query \$query: response was \$response	'query' - The query address.
	Information. Sent when an error occurred while processing a response from SenderBase.	'response' - Raw data of response received.
SMTPAUTH.FWD_SERVER_FAILED_ALERT	SMTP Auth: could not reach forwarding server \$ip with reason: \$why	'ip' - The IP of the remote server.
	Warning. Sent when the SMTP Authentication forwarding server is unreachable.	'why' - Why the error happened.
SMTPAUTH.LDAP_QUERY_FAILED	SMTP Auth: LDAP query failed, see LDAP debug logs for details.	
	Warning. Sent when an LDAP query fails.	
SYSTEM.HERMES_SHUTDOWN_FAILURE. REBOOT	While preparing to \${what}, failed to stop mail server gracefully: \${error}\$what:=reboot	'error' - The error that happened.
	Warning. Sent when there was a problem shutting down the system on reboot.	
SYSTEM.HERMES_SHUTDOWN_FAILURE. SHUTDOWN	While preparing to \${what}, failed to stop mail server gracefully: \${error}\$what:=shut down	'error' - The error that happened.
	Warning. Sent when there was a problem shutting down the system.	

Table 33-7 Listing of Possible System Alerts (continued)

Component/Alert Name	Message and Description	Parameters
SYSTEM.RCPTVALIDATION.UPDATE_FAILED	Error updating recipient validation data: \$why	'why' - The error message.
	Critical. Sent when a recipient validation update failed.	
SYSTEM.SERVICE_TUNNEL.DISABLED	Tech support: Service tunnel has been disabled	
	Information. Sent when a tunnel created for Cisco Support Services is disabled.	
SYSTEM.SERVICE_TUNNEL.ENABLED	Tech support: Service tunnel has been enabled, port \$port	'port' - The port used for the service tunnel.
	Information. Sent when a tunnel created for Cisco Support Services is enabled.	
IPBLOCKD.HOST_ADDED_TO_WHITELIST	The host at \$ip has been added to the blacklist because of an SSH DOS attack.	'ip' - IP address from which a login attempt occurred.
IPBLOCKD.HOST_ADDED_TO_BLACKLIST	The host at \$ip has been permanently added to the ssh whitelist.	
IPBLOCKD.HOST_REMOVED_FROM_BLACKLIST	The host at \$ip has been removed from the blacklist	
	Warning. IP addresses that try to connect to the appliance over SSH but do not provide valid credentials are added to the SSH blacklist if more than 10 failed attempts occur within two minutes. When a user logs in successfully from the same IP address, that IP address is added to the whitelist. Addresses on the whitelist are allowed access even if they are also on the blacklist. Entries are automatically removed from the blacklist after about a day.	

Updater Alerts

Table 33-8 contains a list of the various Updater alerts that can be generated by AsyncOS.

Table 33-8 Listing of Possible Updater Alerts

Alert Name	Message and Description	Parameters
UPDATER.APP.UPDATE_ABANDONED	\$app abandoning updates until a new version is published. The \$app application tried and failed \$attempts times to successfully complete an update. This may be due to a network configuration issue or temporary outage	'app' - The application name. 'attempts' - The number of attempts tried.
	Warning. The application is abandoning the update.	
UPDATER.UPDATERD.MANIFEST_FAILED_ALERT	The updater has been unable to communicate with the update server for at least \$threshold.	'threshold' - Human readable threshold string.
	Warning. Failed to acquire a server manifest.	

Table 33-8 Listing of Possible Updater Alerts (continued)

Alert Name	Message and Description	Parameters
UPDATER.UPDATERD.RELEASE_NOTIFICATION	\$mail_text	'mail_text' - The notification text.
	Warning. Release notification.	'notification_subject' - The notification text.
UPDATER.UPDATERD.UPDATE_FAILED	Unknown error occurred: \$traceback	'traceback' - The traceback.
	Critical. Failed to run an update.	

Outbreak Filter Alerts

Table 33-9 contains a list of the various Outbreak Filter alerts that can be generated by AsyncOS, including a description of the alert and the alert severity. Please note that Outbreak Filters can also be referenced in system alerts for quarantines (the Outbreak quarantine, specifically).

Table 33-9 Listing of Possible Outbreak Filter Alerts

Alert Name	Message and Description	Parameters
VOFGTL_THRESHOLD_ALERT	Outbreak Filters Rule Update Alert:\$text All rules last updated at: \$time on \$date.	'text' - Update alert text.
	Information. Sent when the Outbreak Filters threshold has changed.	'time' - Time of last update. 'date' - Date of last update.
AS.UPDATE_FAILURE	\$engine update unsuccessful. This may be due to transient network or DNS issues, HTTP proxy configuration causing update transmission errors or unavailability of downloads.ironport.com. The specific error on the appliance for this failure is: \$error	'engine' - The engine that failed to update. 'error' - The error that happened.
	Warning. Sent when the anti-spam engine or CASE rules fail to update.	

Clustering Alerts

Table 33-10 contains a list of the various clustering alerts that can be generated by AsyncOS, including a description of the alert and the alert severity.

Table 33-10 Listing of Possible Clustering Alerts

Alert Name	Message and Description	Parameters
CLUSTER.CC_ERROR.AUTH_ERROR	Error connecting to cluster machine \$name at IP \$ip - \$error - \$why\$error:=Machine does not appear to be in the cluster	'name' - The hostname and/or serial number of the machine.
	Critical. Sent when there was an authentication error. This can occur if a machine is not a member of the cluster.	'ip' - The IP of the remote host. 'why' - Detailed text about the error.

Table 33-10 Listing of Possible Clustering Alerts (continued)

Alert Name	Message and Description	Parameters
CLUSTER.CC_ERROR.DROPPED	Error connecting to cluster machine \$name at IP \$ip - \$error - \$why\$error:=Existing connection dropped	'name' - The hostname and/or serial number of the machine.
	Warning. Sent when the connection to the cluster was dropped.	'ip' - The IP of the remote host. 'why' - Detailed text about the error.
CLUSTER.CC_ERROR.FAILED	Error connecting to cluster machine \$name at IP \$ip - \$error - \$why\$error:=Connection failure	'name' - The hostname and/or serial number of the machine.
	Warning. Sent when the connection to the cluster failed.	'ip' - The IP of the remote host. 'why' - Detailed text about the error.
CLUSTER.CC_ERROR.FORWARD_FAILED	Error connecting to cluster machine \$name at IP \$ip - \$error - \$why\$error:=Message forward failed, no upstream connection	'name' - The hostname and/or serial number of the machine.
	Critical. Sent when the appliance was unable to forward data to a machine in the cluster.	'ip' - The IP of the remote host. 'why' - Detailed text about the error.
CLUSTER.CC_ERROR.NOROUTE	Error connecting to cluster machine \$name at IP \$ip - \$error - \$why\$error:=No route found	'name' - The hostname and/or serial number of the machine.
	Critical. Sent when the machine was unable to obtain a route to another machine in the cluster.	'ip' - The IP of the remote host. 'why' - Detailed text about the error.
CLUSTER.CC_ERROR.SSH_KEY	Error connecting to cluster machine \$name at IP \$ip - \$error - \$why\$error:=Invalid host key	'name' - The hostname and/or serial number of the machine.
	Critical. Sent when there was an invalid SSH host key.	'ip' - The IP of the remote host. 'why' - Detailed text about the error.

Table 33-10 Listing of Possible Clustering Alerts (continued)

Alert Name	Message and Description	Parameters
CLUSTER.CC_ERROR.TIMEOUT	Error connecting to cluster machine \$name at IP \$ip - \$error - \$why\$error:=Operation timed out	'name' - The hostname and/or serial number of the machine.
	Warning. Sent when the specified operation timed out.	'ip' - The IP of the remote host. 'why' - Detailed text about the error.
CLUSTER.CC_ERROR_NOIP	Error connecting to cluster machine \$name - \$error - \$why	'name' - The hostname and/or serial number of the machine.
	Critical. Sent when the appliance could not obtain a valid IP address for another machine in the cluster.	'why' - Detailed text about the error.
CLUSTER.CC_ERROR_NOIP.AUTH_ERROR	Error connecting to cluster machine \$name - \$error - \$why\$error:=Machine does not appear to be in the cluster	'name' - The hostname and/or serial number of the machine.
	Critical. Sent when there was an authentication error connecting to a machine in a cluster. This can occur if a machine is not a member of the cluster.	'why' - Detailed text about the error.
CLUSTER.CC_ERROR_NOIP.DROPPED	Error connecting to cluster machine \$name - \$error - \$why\$error:=Existing connection dropped	'name' - The hostname and/or serial number of the machine.
	Warning. Sent when the machine was unable to obtain a valid IP address for another machine in the cluster and the connection to the cluster was dropped.	'why' - Detailed text about the error.
CLUSTER.CC_ERROR_NOIP.FAILED	Error connecting to cluster machine \$name - \$error - \$why\$error:=Connection failure	'name' - The hostname and/or serial number of the machine.
	Warning. Sent when there was an unknown connection failure and the machine was unable to obtain a valid IP address for another machine in the cluster.	'why' - Detailed text about the error.
CLUSTER.CC_ERROR_NOIP.FORWARD_FAILED	Error connecting to cluster machine \$name - \$error - \$why\$error:=Message forward failed, no upstream connection	'name' - The hostname and/or serial number of the machine.
	Critical. Sent when the machine was unable to obtain a valid IP address for another machine in the cluster and the appliance was unable to forward data to the machine.	'why' - Detailed text about the error.
CLUSTER.CC_ERROR_NOIP.NOROUTE	Error connecting to cluster machine \$name - \$error - \$why\$error:=No route found	'name' - The hostname and/or serial number of the machine.
	Critical. Sent when the machine was unable to obtain a valid IP address for another machine in the cluster and it was unable to obtain a route to the machine.	'why' - Detailed text about the error.

Table 33-10 Listing of Possible Clustering Alerts (continued)

Alert Name	Message and Description	Parameters
CLUSTER.CC_ERROR_NOIP.SSH_KEY	Error connecting to cluster machine \$name - \$error - \$why\$error:=Invalid host key	'name' - The hostname and/or serial number of the machine.
	Critical. Sent when the machine was unable to obtain a valid IP address for another machine in the cluster and was unable to obtain a valid SSH host key.	'why' - Detailed text about the error.
CLUSTER.CC_ERROR_NOIP.TIMEOUT	Error connecting to cluster machine \$name - \$error - \$why\$error:=Operation timed out	'name' - The hostname and/or serial number of the machine.
	Warning. Sent when the machine was unable to obtain a valid IP address for another machine in the cluster and the specified operation timed out.	'why' - Detailed text about the error.
CLUSTER.SYNC.PUSH_ALERT	Overwriting \$sections on machine \$name	'name' - The hostname and/or serial number of the machine.
	Critical. Sent when configuration data has gotten out of sync and has been sent to a remote host.	'sections' - List of cluster sections being sent.

Changing Network Settings

This section describes the features used to configure the network operation of the appliance. These features give you direct access to the hostname, DNS, and routing settings that you configured via the System Setup Wizard (or the `systemsetup` command) in [Using the System Setup Wizard, page 3-13](#).

The following features are described:

- `sethostname`
- DNS Configuration (GUI and via the `dnsconfig` command)
- Routing Configuration (GUI and via the `routeconfig` and `setgateway` commands)
- `dnsflush`
- Password
- Network Access
- Login Banner

Changing the System Hostname

The hostname is used to identify the system at the CLI prompt. You must enter a fully-qualified hostname. The `sethostname` command sets the name of the appliance. The new hostname does not take effect until you issue the `commit` command.

The sethostname Command

```
oldname.example.com> sethostname
```

```
[oldname.example.com]> mail3.example.com
```

```
oldname.example.com>
```

For the hostname change to take effect, you must enter the `commit` command. After you have successfully committed the hostname change, the new name appears in the CLI prompt:

```
oldname.example.com> commit
```

```
Please enter some comments describing your changes:
```

```
[> Changed System Hostname
```

```
Do you want to save the current configuration for rollback? [Y]> n
```

```
Changes committed: Fri May 23 11:42:12 2014 GMT
```

The new hostname appears in the prompt as follows: `mail3.example.com>`

Configuring Domain Name System (DNS) Settings

You can configure the DNS settings for your appliance through the DNS page on the Network menu of the GUI, or via the `dnsconfig` command.

You can configure the following settings:

- whether to use the Internet's DNS servers or your own, and which specific server(s) to use
- which interface to use for DNS traffic
- the number of seconds to wait before timing out a reverse DNS lookup
- clear DNS cache

Specifying DNS Servers

AsyncOS can use the Internet root DNS servers, your own DNS servers, or the Internet root DNS servers and authoritative DNS servers you specify. When using the Internet root servers, you may specify alternate servers to use for specific domains. Since an alternate DNS server applies to a single domain, it must be authoritative (provide definitive DNS records) for that domain.

AsyncOS supports “splitting” DNS servers when not using the Internet’s DNS servers. If you are using your own internal server, you can also specify exception domains and associated DNS servers.

When setting up “split DNS,” you should set up the in-addr.arpa (PTR) entries as well. So, for example, if you want to redirect “.eng” queries to the nameserver 1.2.3.4 and all the .eng entries are in the 172.16 network, then you should specify “eng,16.172.in-addr.arpa” as the domains in the split DNS configuration.

Multiple Entries and Priority

For each DNS server you enter, you can specify a numeric priority. AsyncOS will attempt to use the DNS server with the priority closest to 0. If that DNS server is not responding AsyncOS will attempt to use the server at the next priority. If you specify multiple entries for DNS servers with the same priority, the system randomizes the list of DNS servers at that priority every time it performs a query. The system then waits a short amount of time for the first query to expire or “time out” and then a slightly longer amount of time for the second, etc. The amount of time depends on the exact total number of DNS servers and priorities that have been configured. The timeout length is the same for all IP addresses at any particular priority. The first priority gets the shortest timeout, each subsequent priority gets a longer timeout. Further, the timeout period is roughly 60 seconds. If you have one priority, the timeout for each server at that priority will be 60 seconds. If you have two priorities, the timeout for each server at the first priority will be 15 seconds, and each server at the second priority will be 45 seconds. For three priorities, the timeouts are 5, 10, 45.

For example, suppose you configure four DNS servers, with two of them at priority 0, one at priority 1, and one at priority 2:

Table 33-11 Example of DNS Servers, Priorities, and Timeout Intervals

Priority	Server(s)	Timeout (seconds)
0	1.2.3.4, 1.2.3.5	5, 5
1	1.2.3.6	10
2	1.2.3.7	45

AsyncOS will randomly choose between the two servers at priority 0. If one of the priority 0 servers is down, the other will be used. If both of the priority 0 servers are down, the priority 1 server (1.2.3.6) is used, and then, finally, the priority 2 (1.2.3.7) server.

The timeout period is the same for both priority 0 servers, longer for the priority 1 server, and longer still for the priority 2 server.

Using the Internet Root Servers

The AsyncOS DNS resolver is designed to accommodate the large number of simultaneous DNS connections required for high-performance email delivery.



Note

If you choose to set the default DNS server to something other than the Internet root servers, that server must be able to recursively resolve queries for domains for which it is not an authoritative server.

Reverse DNS Lookup Timeout

The appliance attempts to perform a “double DNS lookup” on all remote hosts connecting to a listener for the purposes of sending or receiving email. [That is: the system acquires and verifies the validity of the remote host's IP address by performing a double DNS lookup. This consists of a reverse DNS (PTR) lookup on the IP address of the connecting host, followed by a forward DNS (A) lookup on the results of the PTR lookup. The system then checks that the results of the A lookup match the results of the PTR lookup. If the results do not match, or if an A record does not exist, the system only uses the IP address to match entries in the Host Access Table (HAT).] This particular timeout period applies only to this lookup and is not related to the general DNS timeout discussed in [Multiple Entries and Priority](#), page 33-55.

The default value is 20 seconds. You can disable the reverse DNS lookup timeout globally across all listeners by entering ‘0’ as the number of seconds.

If the value is set to 0 seconds, the reverse DNS lookup is not attempted, and instead the standard timeout response is returned immediately. This also prevents the appliance from delivering mail to domains that require TLS-verified connections if the receiving host’s certificate has a common name (CN) that maps to the host’s IP lookup.

DNS Alert

Occasionally, an alert may be generated with the message “Failed to bootstrap the DNS cache” when an appliance is rebooted. The message means that the system was unable to contact its primary DNS servers, which can happen at boot time if the DNS subsystem comes online before network connectivity is established. If this message appears at other times, it could indicate network issues or that the DNS configuration is not pointing to a valid server.

Clearing the DNS Cache

The Clear Cache button from the GUI, or the `dnsflush` command (for more information about the `dnsflush` command, see the *Cisco AsyncOS CLI Reference Guide*), clears all information in the DNS cache. You may choose to use this feature when changes have been made to your local DNS system. The command takes place immediately and may cause a temporary performance degradation while the cache is repopulated.

Configuring DNS Settings via the Graphical User Interface

Procedure

-
Step 1	Select Network > DNS .
Step 2	Click Edit Settings .
Step 3	Select whether to use the Internet’s root DNS servers or your own internal DNS server or the Internet’s root DNS servers and specify alternate DNS servers.
Step 4	If you want to use your own DNS server(s) enter the server ID and click Add Row . Repeat this for each server. When entering your own DNS servers, specify a priority as well. For more information, see Specifying DNS Servers , page 33-54.
Step 5	If you want to specify alternate DNS servers for certain domains, enter the domain and the alternate DNS server IP address. Click Add Row to add additional domains.

**Note**

You can enter multiple domains for a single DNS server by using commas to separate domain names. You can also enter multiple DNS servers by using commas to separate IP addresses.

- Step 6** Choose an interface for DNS traffic.
- Step 7** Enter the number of seconds to wait before cancelling a reverse DNS lookup.
- Step 8** You can also clear the DNS cache by clicking **Clear Cache**.
- Step 9** Submit and commit your changes.

Configuring TCP/IP Traffic Routes

Some network environments require the use of traffic routes other than the standard default gateway.

The Email Security appliance can use both Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) static routes.

You can manage static routes via the CLI, using the `routeconfig` command, or use the following procedure.

Procedure

- Step 1** Select **Network > Routing**.
- Step 2** Click **Add Route** for the type of static route you want to create (IPv4 or IPv6).
- Step 3** Enter a name for the route.
- Step 4** Enter the destination IP address.
- Step 5** Enter the Gateway IP address.
- Step 6** Submit and commit your changes.

Configuring the Default Gateway

You can configure the default gateway using the `setgateway` command in the CLI or use the following procedure.

Procedure

- Step 1** Select **Network > Routing**.
- Step 2** Click **Default Route** in the route listing for the Internet Protocol version you want to modify.
- Step 3** Change the Gateway IP address.
- Step 4** Submit and commit your changes.

Configuring SSL Settings

You can configure the SSL settings for the appliance using SSL Configuration Settings page or `sslconfig` command.

Procedure

-
- Step 1** Click **System Administration > SSL Configuration Settings**.
- Step 2** Click **Edit Settings**.
- Step 3** Depending on your requirements, do the following:
- Set GUI HTTPS SSL settings. Under GUI HTTPS, specify the SSL methods and ciphers that you want to use.
 - Set Inbound SMTP SSL settings. Under Inbound SMTP, specify the SSL methods and ciphers that you want to use.
 - Set Outbound SMTP SSL settings. Under Outbound SMTP, specify the SSL methods and ciphers that you want to use.



Note You cannot enable SSL v2 and TLS v1 methods simultaneously. However, you can enable these methods in conjunction with SSL v3 method.

- Step 4** Click **Submit**.
- Step 5** Click **Commit Changes**.
-

Disabling SSLv3 for Enhanced Security

For enhanced security, you can disable SSLv3 for the following services:

- Updater
- URL Filtering
- End User Quarantine
- LDAP

Use the `ssl3config` command in CLI to enable or disable SSLv3 for the above services. The following example shows how to disable SSLv3 for End User Quarantine.

```
mail.example.com> ssl3config
```

```
Current SSLv3 Settings:
```

```
-----
UPDATER      : Enabled
WEBSECURITY  : Enabled
EUQ          : Enabled
LDAP         : Enabled
-----
```

```
Choose the operation you want to perform:
```

```
- SETUP - Toggle SSLv3 settings.
```

```
[> setup
```

```
Choose the service to toggle SSLv3 settings:
1. EUQ Service
2. LDAP Service
3. Updater Service
4. Web Security Service
[1]>

Do you want to enable SSLv3 for EUQ Service ? [Y]>n

Choose the operation you want to perform:
- SETUP - Toggle SSLv3 settings.
[]>
```

System Time

To set the System Time on your appliance, set the Time Zone used, or select an NTP server and query interface, use the Time Zone or Time Settings page from the System Administration menu in the GUI or use the following commands in the CLI: `ntpconfig`, `settime`, and `settz`.

You can also verify the time zone files used by AsyncOS on the **System Administration > Time Settings** page or using the `tzupdate` CLI command.

Selecting a Time Zone

The Time Zone page (available via the System Administration menu in the GUI) displays the time zone for your appliance. You can select a specific time zone or GMT offset.

Procedure

-
Step 1	Click Edit Settings on the System Administration > Time Zone page.
Step 2	Select a Region, country, and time zone from the pull-down menus.
Step 3	Submit and commit your changes.
-

Selecting a GMT Offset

Procedure

-
Step 1	Click Edit Settings on the System Administration > Time Zone page.
Step 2	Select GMT Offset from the list of regions.
Step 3	Select an offset in the Time Zone list. The offset refers to the amount of hours that must be added/subtracted in order to reach GMT (the Prime Meridian). Hours preceded by a minus sign (“-”) are east of the Prime Meridian. A plus sign (“+”) indicates west of the Prime Meridian.
Step 4	Submit and commit your changes.
-

Editing Time Settings

You can edit the time settings for the appliance using one of the following methods:

- Using the Networking Time Protocol (NTP)
- Manually

Setting Appliance System Time Using the Network Time Protocol (NTP)

Procedure

- Step 1** Navigate to the System Administration > Time Settings page.
- Step 2** Click **Edit Settings**.
- Step 3** In the Time Keeping Method section, select Use Network Time Protocol.
- Step 4** Enter an NTP server address and click **Add Row**. You can add multiple NTP servers.
- Step 5** To delete an NTP server from the list, click the trash can icon for that server.
- Step 6** Select an interface for NTP queries. This is the IP address from which NTP queries should originate.
- Step 7** Submit and commit your changes.
-

Setting Appliance System Time Manually

Procedure

- Step 1** Navigate to the System Administration > Time Settings page.
- Step 2** Click **Edit Settings**.
- Step 3** In the Time Keeping Method section, select Set Time Manually.
- Step 4** Enter the month, day, year, hour, minutes, and seconds.
- Step 5** Select A.M or P.M.
- Step 6** Submit and commit your changes.
-

Customizing Your View

- [Using Favorite Pages, page 33-60](#)
- [Setting User Preferences, page 33-61](#)

Using Favorite Pages

(Locally-authenticated administrative users only.) You can create a quick-access list of the pages you use most.

To	Do This
Add pages to your favorites list	Navigate to the page to add, then choose Add This Page To My Favorites from the My Favorites menu near the top right corner of the window. No commit is necessary for changes to My Favorites.
Reorder favorites	Choose My Favorites > View All My Favorites and drag favorites into the desired order.
Delete favorites	Choose My Favorites > View All My Favorites and delete favorites.
Go to a favorite page	Choose a page from the My Favorites menu near the top right corner of the window.
View or build a custom reporting page	See My Reports Page, page 28-5 .

Setting User Preferences

Local users can define preference settings, such as language, specific to each account. These settings apply by default when the user first logs into the appliance. The preference settings are stored for each user and are the same regardless from which client machine the user logs into the appliance.

When users change these settings but do not commit the changes, the settings revert to the default values when they log in again.



Note

This feature is not available to externally-authenticated users. These users can choose a language directly from the Options menu.

Procedure

- Step 1** Log into the appliance with the user account for which you want to define preference settings.
- Step 2** Choose **Options > Preferences**. The options menu is at the top right side of the window.
- Step 3** Click **Edit Preferences**.
- Step 4** Configure settings:

Preference Setting	Description
Language Display	The language AsyncOS for Web uses in the web interface and CLI.
Landing Page	The page that displays when the user logs into the appliance.
Reporting Time Range Displayed (default)	The default time range that displays for reports on the Reporting tab.
Number of Reporting Rows Displayed	The number of rows of data shown for each report by default.

- Step 5** Submit and commit your changes.

- Step 6** Click the **Return to previous page** link at the bottom of the page.
-

Overriding Internet Explorer Compatibility Mode

For better web interface rendering, Cisco recommends that you enable Internet Explorer Compatibility Mode Override.



Note

If enabling this feature is against your organizational policy, you may disable this feature.

Procedure

- Step 1** Click System **Administration** > **General Settings**.
- Step 2** Select **Override IE Compatibility Mode** check box.
- Step 3** Submit and commit your changes.
-



Managing and Monitoring Using the CLI

- [Overview of Managing and Monitoring Using the CLI, page 34-1](#)
- [Monitoring Using the CLI, page 34-6](#)
- [Managing the Email Queue, page 34-22](#)
- [SNMP Monitoring, page 34-36](#)

Overview of Managing and Monitoring Using the CLI

Managing and monitoring the Email Security appliance using the CLI includes these types of tasks:

- Monitoring message activity.
 - The raw number of messages, recipients, and bounce recipients that the appliance is processing in the email pipeline
 - The hourly rate of message delivery or message bounces based on the last one-minute, five-minute, or fifteen-minute period
- Monitoring system resources. Examples:
 - Memory usage
 - Disk space
 - Number of connections
- Monitoring possible system disfunction using the Simple Network Management Protocol (SNMP). Examples:
 - Fan failure
 - Update failure
 - Abnormally high appliance temperature
- Managing email within the pipeline. Examples:
 - Deleting recipients in the queue
 - Redirecting messages to another host
 - Clear the queue by deleting recipients or redirecting the messages
 - Suspend or resume email receiving, delivery, or work queue processing
 - Locate specific messages

Reading the Available Components of Monitoring

- [Reading the Available Components of Monitoring, page 34-2](#)
- [Reading the Event Counters, page 34-2](#)
- [Reading the System Gauges, page 34-4](#)
- [Reading the Rates of Delivered and Bounced Messages, page 34-6](#)

Reading the Event Counters

Counters provide a running total of various events in the system. For each counter, you can view the total number of events that have occurred since the counter was reset, since the last system reboot, and over the system's lifetime.

Counters increment each time an event occurs and are displayed in three versions:

Reset	Since the last counter reset with the <code>resetcounters</code> command
Uptime	Since the last system reboot
Lifetime	Total through the lifetime of the Cisco appliance

Table 34-1 lists the available counters and their description when monitoring the Cisco appliance.



Note

This is the entire list. The displayed counters vary depending on which display option or command you choose. Use this list as a reference.

Table 34-1 **Counters**

Statistic	Description
Receiving	
Messages Received	Messages received into the delivery queue.
Recipients Received	Recipients on all received messages.
Generated Bounce Recipients	Recipients for which bounces have been generated by the system and inserted into the delivery queue.

Table 34-1 Counters (continued)

Statistic	Description
Rejection	
Rejected Recipients	Recipients that have been denied receiving into the delivery queue due to the Recipient Access Table (RAT), or unexpected protocol negotiation including premature connection termination.
Dropped Messages	Messages that have been denied receiving into the delivery queue due to a filter drop action match or have been received by a Black Hole queuing listener. Messages directed to <code>/dev/null</code> entries in the alias table also are considered dropped messages. Messages dropped by anti-spam filtering (if it has been enabled on the system) also increment this counter.
Queue	
Soft Bounced Events	Number of soft bounce events — a message that soft bounces multiple times has multiple soft bounce events.
Completion	
Completed Recipients	Total of all hard bounced recipients, delivered recipients, and deleted recipients. Any recipient that is removed from the delivery queue.
Hard Bounced Recipients	Total of all DNS hard bounces, 5XX hard bounces, filter hard bounces, expired hard bounces and other hard bounces. A failed attempt to deliver message to a recipient that results in immediate termination of that delivery.
DNS Hard Bounces	DNS error encountered while trying to deliver a message to a recipient.
5XX Hard Bounces	The destination mail server returned a “5XX” response code while trying to deliver a message to a recipient.
Expired Hard Bounces	Message recipients that have exceeded the maximum time allowed in the delivery queue or the maximum number of connection attempts.
Filter Hard Bounces	Recipient delivery has been preempted by a matching filter bounce action. Messages dropped by anti-spam filtering (if it has been enabled on the system) also increment this counter.
Other Hard Bounces	An unexpected error during message delivery or a message recipient was explicitly bounced via the <code>bouncerecipients</code> command.
Delivered Recipients	Message successfully delivered to a recipient.
Deleted Recipients	Total of message recipients explicitly deleted via the <code>deleterecipients</code> command or was a Global Unsubscribe Hit.
Global Unsubscribe Hits	Message recipient was deleted due to a matching global unsubscribe setting.
Current IDs	
Message ID (MID)	The last Message ID to have been assigned to a message inserted into the delivery queue. A MID is associated with every message received by the Cisco appliance and can be tracked in mail logs. The MID resets to zero at 2^{31} .

Table 34-1 Counters (continued)

Statistic	Description
Injection Connection ID (ICID)	The last Injection Connection ID to have been assigned to a connection to a listener interface. The ICID rolls over (resets to zero) at 2^{31} .
Delivery Connection ID (DCID)	The last Delivery Connection ID to have been assigned to a connection to a destination mail server. The DCID rolls over (resets to zero) at 2^{31} .

Reading the System Gauges

Gauges show the current utilization of a system resource such as memory, disk space, or active connections.

Table 34-2 lists the available gauges and their description when monitoring the Cisco appliance.


Note

This is the entire list. The displayed gauges will vary depending upon which display option or command you choose. Use this list as a reference.

Table 34-2 Gauges

Statistic	Description
System Gauges	
RAM Utilization	Percentage of physical RAM (Random Access Memory) being used by the system.
CPU Utilization	Percentage of CPU usage.
Disk I/O Utilization	Percentage of Disk I/O being used. Note The Disk I/O Utilization gauge does not display a reading against a scale of a known value. Rather, it displays the I/O utilization the system has seen thus far and scales against the maximum value since the last reboot. So, if the gauge displays 100%, the system is experiencing the highest level of I/O utilization seen since boot (which may not necessarily represent 100% of the physical Disk I/O of the entire system).
Resource Conservation	A value between 0 and 60 or 999. Numbers from 0 to 60 represent the degree to which the system is decreasing its acceptance of messages in order to prevent the rapid depletion of critical system resources. Higher numbers represent a higher degree of decreased acceptance. Zero represents no decrease in acceptance. If this gauge displays 999, the system has entered “Resource Conservation mode,” and it will accept no messages. Alert messages are sent whenever the system enters or exits Resource Conservation mode.
Disk Utilization: Logs	Percentage of disk being used for logs, displayed as <code>LogUsed</code> in the status logs and <code>log_used</code> in the XML status.

Table 34-2 Gauges (continued)

Statistic	Description
Connections Gauges	
Current Inbound Connections	Current inbound connections to the listener interfaces.
Current Outbound Connections	Current outbound connections to destination mail servers.
Queue Gauges	
Active Recipients	Message recipients in the delivery queue. Total of Unattempted Recipients and Attempted Recipients.
Unattempted Recipients	A subcategory of Active Recipients. Message recipients in queue for which delivery has not yet been attempted.
Attempted Recipients	A subcategory of Active Recipients. Message recipients in queue for which delivery has been attempted but failed due to a Soft Bounces Event.
Messages in Work Queue	The number of messages waiting to be processed by alias table expansion, masquerading, anti-spam, anti-virus scanning, message filters, and LDAP queries prior to being enqueued.
Messages in Quarantine	The unique number of messages in any quarantine, plus messages that have been released or deleted but not yet acted upon. For example, if you release all quarantined messages from Outbreak, the total messages for Outbreak would become zero immediately, but this field still reflects the quarantined messages until they were all delivered.
Destinations in Memory	The number of destinations domains in memory. For each domain with a message destined to be delivered, a destination object is created in memory. After all the mail for that domain has been delivered, the destination object is retained for another 3 hours. After 3 hours, if no new messages are bound for that domain, the object is expired so that the destination is no longer reported (for example, in the <code>tophosts</code> command). If you are delivering mail only to one domain, this counter will be “1.” If you have never received or sent any messages (or no messages have been processed by the appliance in many hours), the counter will be “0.” If you are using Virtual Gateways, destination domains for each Virtual Gateway will have a separate destination object. (For example, yahoo.com will count as 3 destination objects if you are delivering to yahoo.com from 3 different Virtual Gateways).
Kilobytes Used	Queue storage used in kilobytes.
Kilobytes in Quarantine	Queue storage used for quarantined messages. The value is calculated as the message size plus 30 bytes for each recipient, totaled for the “Messages in Quarantine” as counted above. Note that this calculation will usually <i>overestimate</i> the space used.
Kilobytes Free	Queue storage remaining in kilobytes.

Reading the Rates of Delivered and Bounced Messages

All rates are shown as the average rate an event occurs per hour at the specific point in time the query is made. Rates are calculated for three intervals, the average rate per hour over the past one (1) minute, the past five (5) minutes, and the past fifteen (15) minutes.

For example, if the Cisco appliance receives 100 recipients in a single minute, then the rate for the 1 minute interval will be 6,000 per hour. The rate for the 5-minute interval will be 1,200 per hour, and the 15-minute rate will be 400 per hour. The rates are calculated to indicate what the average rate for the hour would be if the rate for the one minute period continued. Therefore, 100 messages each minute would yield a higher rate than 100 messages over 15 minutes.

Table 34-3 lists the available rates and their description when monitoring the Cisco appliance.



Note

This is the entire list. The displayed rates will vary depending upon which display option or command you choose. Use this list as a reference.

Table 34-3 Rates

Statistic	Description
Messages Received	Rate of messages inserted into the delivery queue per hour.
Recipients Received	Rate of the number of recipients on all messages inserted into the delivery queue per hour.
Soft Bounced Events	Rate of the number of soft bounce events per hour. (A message that soft bounces multiple times has multiple soft bounce events.)
Completed Recipients	Rate of the total of all hard bounced recipients, delivered recipients and deleted recipients. Any recipient that is removed from the delivery queue is considered completed.
Hard Bounced Recipients	Rate of the total of all DNS hard bounces, 5XX hard bounces, filter hard bounces, expired hard bounces and other hard bounces per hour. A failed attempt to deliver a message to a recipient that results in immediate termination of that delivery is a hard bounce.
Delivered Recipients	Rate of messages successfully delivered to a recipient per hour.

Monitoring Using the CLI

- [Monitoring the Email Status, page 34-7](#)
- [Monitoring Detailed Email Status, page 34-8](#)
- [Monitoring the Status of a Mail Host, page 34-11](#)
- [Determining the Make-up of the Email Queue, page 34-15](#)
- [Displaying Real-time Activity, page 34-16](#)
- [Monitoring Inbound Email Connections, page 34-19](#)
- [Checking the DNS Status, page 34-20](#)
- [Resetting Email Monitoring Counters, page 34-21](#)
- [Identifying Active TCP/IP Services, page 34-22](#)

Monitoring the Email Status

You may want to monitor the status of email operations on the Cisco appliance. The `status` command returns a subset of the monitored information about email operations. The statistics returned displayed in one of two fashions: counters and gauges. Counters provide a running total of various events in the system. For each counter, you can view the total number of events that have occurred since the counter was reset, since the last system reboot, and over the system's lifetime. Gauges show the current utilization of a system resource such as memory, disk space, or active connections.

For a description of each item, see [Overview of Managing and Monitoring Using the CLI, page 34-1](#).

Table 34-4 Mail Status

Statistic	Description
Status as of	Displays the current system time and date.
Last counter reset	Displays the last time the counters were reset.
System status	Online, offline, receiving suspended, or delivery suspended. Note that the status will be “receiving suspended” only when <i>all</i> listeners are suspended. The status will be “offline” when receiving and delivery are suspended for <i>all</i> listeners.
Oldest Message	Displays the oldest message waiting to be delivered by the system.
Features	Displays any special features installed on the system by the <code>featurekey</code> command.

Example

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

```
Status as of: Thu Oct 21 14:33:27 2004 PDT
Up since: Wed Oct 20 15:47:58 2004 PDT (22h 45m 29s)
Last counter reset: Never
System status: Online
Oldest Message: 4 weeks 46 mins 53 secs
```

```
Counters: Reset Uptime Lifetime
Receiving
Messages Received 62,049,822 290,920 62,049,822
Recipients Received 62,049,823 290,920 62,049,823
Rejection
Rejected Recipients 3,949,663 11,921 3,949,663
```

```

Dropped Messages      11,606,037      219      11,606,037

Queue

Soft Bounced Events   2,334,552      13,598      2,334,552

Completion

Completed Recipients   50,441,741      332,625      50,441,741

Current IDs

Message ID (MID)                                99524480

Injection Conn. ID (ICID)                        51180368

Delivery Conn. ID (DCID)                        17550674


Gauges:                                Current

Connections

Current Inbound Conn.      0

Current Outbound Conn.     14

Queue

Active Recipients          7,166

Messages In Work Queue      0

Messages In Quarantine     16,248

Kilobytes Used              387,143

Kilobytes In Quarantine     338,206

Kilobytes Free              39,458,745


mail3.example.com>

```

Monitoring Detailed Email Status

The `status detail` command returns complete monitored information about email operations. The statistics returned are displayed in one of three categories: counters, rates, and gauges. Counters provide a running total of various events in the system. For each counter, you can view the total number of events that have occurred since the counter was reset, since the last system reboot, and over the system's lifetime. Gauges show the current utilization of a system resource such as memory, disk space, or active connections. All rates are shown as the average rate an event occurs per hour at the specific point in time.

the query is made. Rates are calculated for three intervals, the average rate per hour over the past one (1) minute, the past five (5) minutes, and the past fifteen (15) minutes. For a description of each item, see [Overview of Managing and Monitoring Using the CLI, page 34-1](#).

Example

```
mail3.example.com> status detail
```

```
Status as of:          Thu Jun 30 13:09:18 2005 PDT

Up since:             Thu Jun 23 22:21:14 2005 PDT (6d 14h 48m 4s)

Last counter reset:   Tue Jun 29 19:30:42 2004 PDT

System status:        Online

Oldest Message:       No Messages

Feature - IronPort Anti-Spam: 17 days

Feature - Sophos:      Dormant/Perpetual

Feature - Outbreak Filters: Dormant/Perpetual

Feature - Central Mgmt: Dormant/Perpetual
```

Counters:	Reset	Uptime	Lifetime
Receiving			
Messages Received	2,571,967	24,760	3,113,176
Recipients Received	2,914,875	25,450	3,468,024
Gen. Bounce Recipients	2,165	0	7,451
Rejection			
Rejected Recipients	1,019,453	792	1,740,603
Dropped Messages	1,209,001	66	1,209,028
Queue			
Soft Bounced Events	11,236	0	11,405
Completion			
Completed Recipients	2,591,740	49,095	3,145,002
Hard Bounced Recipients	2,469	0	7,875

DNS Hard Bounces	199	0	3,235
5XX Hard Bounces	2,151	0	4,520
Expired Hard Bounces	119	0	120
Filter Hard Bounces	0	0	0
Other Hard Bounces	0	0	0
Delivered Recipients	2,589,270	49,095	3,137,126
Deleted Recipients	1	0	1
Global Unsub. Hits	0	0	0
DomainKeys Signed Msgs	10	9	10

Current IDs

Message ID (MID)	7615199
Injection Conn. ID (ICID)	3263654
Delivery Conn. ID (DCID)	1988479

Rates (Events Per Hour):

	1-Minute	5-Minutes	15-Minutes
Receiving			
Messages Received	180	300	188
Recipients Received	180	300	188
Queue			
Soft Bounced Events	0	0	0
Completion			
Completed Recipients	360	600	368
Hard Bounced Recipients	0	0	0
Delivered Recipients	360	600	368

Gauges:

	Current
System	
RAM Utilization	1%

```

CPU Utilization

    MGA                                0%

    AntiSpam                          0%

    AntiVirus                         0%

Disk I/O Utilization                0%

Resource Conservation                0

Connections

    Current Inbound Conn.             0

    Current Outbound Conn.           0

Queue

    Active Recipients                 0

    Unattempted Recipients            0

    Attempted Recipients              0

Messages In Work Queue              0

Messages In Quarantine              19

Destinations In Memory              3

Kilobytes Used                      473

    Kilobytes In Quarantine           473

Kilobytes Free                      39,845,415

```

**Note**

A case could exist in a newly installed appliance where the oldest message counter shows a message but, in fact, there are no recipients shown in counters. If the remote host is connecting and in the process of receiving a message very slowly (that is, it takes minutes to receive a message), you might see that the recipients received counter displays “0” but the oldest message counter displays “1.” This is because the oldest message counter displays messages in progress. The counter will be reset if the connection is eventually dropped.

Monitoring the Status of a Mail Host

If you suspect delivery problems to a specific recipient host or you want to gather information on a Virtual Gateway address, the `hoststatus` command displays this information. The `hoststatus` command returns monitoring information about email operations relating to a specific recipient host. The command requires that you enter the domain of the host information to be returned. DNS information stored in the AsyncOS cache and the last error returned from the recipient host is also given.

Data returned is cumulative since the last `resetcounters` command. The statistics returned are displayed in two categories: counters and gauges. For a description of each item, see [Overview of Managing and Monitoring Using the CLI, page 34-1](#).

In addition, these other data are returned specific to the `hoststatus` command.

Table 34-5 Additional Data in the `hoststatus` Command

Statistic	Description
Pending Outbound Connections	Pending, or “embryonic” connections to the destination mail host, as opposed to open and working connections. Pending Outbound Connections are connections which have not yet gotten to the protocol greeting stage.
Oldest Message	The age of the oldest active recipient in the delivery queue for this domains. This counter is useful for determining the age of a message in the queue that cannot be delivered because of soft bounce events and/or a downed host.
Last Activity	This field is updated each time a message delivery is attempted to that host.
Ordered IP Addresses	This field contains the TTL (time to live) for IP addresses, their preference according to MX records, and the actual addresses. An MX record designates the mail server IP address for a domain. A domain may have multiple MX records. Each MX record mail server is assigned a priority. The MX record with the lowest priority number is given preference.
Last 5XX error	This field contains the most recent “5XX” status code and description returned by the host. This is only displayed if there is an 5XX error.
MX Records	An MX record designates the mail server IP address for a domain. A domain may have multiple MX records. Each MX record mail server is assigned a priority. The MX record with the lowest priority number is given preference.
SMTP Routes for this host	If SMTP routes are defined for this domain, they are listed here.
Last TLS Error	This field contains a description of the the most recent outgoing TLS connection error and the type of TLS connection that the appliance tried to establish. This is only displayed if there is a TLS error.

Virtual Gateway

The following Virtual Gateway information is only displayed if you have set up Virtual Gateway addresses (see [Configuring the Gateway to Receive Email](#).)

Table 34-6 Additional Virtual Gateway Data in the `hoststatus` Command

Statistic	Description
Host up/down	Same definition as global <code>hoststatus</code> field of the same name — tracked per Virtual Gateway address.
Last Activity	Same definition as global <code>hoststatus</code> field of the same name — tracked per Virtual Gateway address.
Recipients	This field also corresponds to the same definition as the global <code>hoststatus</code> command. Active Recipients field — tracked per Virtual Gateway address.
Last 5XX error	This field contains the most recent 5XX status code and description returned by the host. This is only displayed if there is a 5XX error.

Example

```
mail3.example.com> hoststatus

Recipient host:

[]> aol.com

Host mail status for: 'aol.com'

Status as of:          Tue Mar 02 15:17:32 2010

Host up/down:          up

Counters:

Queue

    Soft Bounced Events                0

Completion

    Completed Recipients                1
    Hard Bounced Recipients            1
    DNS Hard Bounces                    0
    5XX Hard Bounces                    1
    Filter Hard Bounces                  0
    Expired Hard Bounces                 0
    Other Hard Bounces                   0
    Delivered Recipients                 0
    Deleted Recipients                   0

Gauges:

Queue

    Active Recipients                    0
    Unattempted Recipients               0
    Attempted Recipients                  0
```


Connections

```

Current Outbound Connections      0

Pending Outbound Connections      0

```

```
Oldest Message      No Messages
```

```
Last Activity      Tue Mar 02 15:17:32 2010
```

```
Ordered IP addresses: (expiring at Tue Mar 02 16:17:32 2010)
```

Preference	IPs
15	64.12.137.121 64.12.138.89 64.12.138.120
15	64.12.137.89 64.12.138.152 152.163.224.122
15	64.12.137.184 64.12.137.89 64.12.136.57
15	64.12.138.57 64.12.136.153 205.188.156.122
15	64.12.138.57 64.12.137.152 64.12.136.89
15	64.12.138.89 205.188.156.154 64.12.138.152
15	64.12.136.121 152.163.224.26 64.12.137.184
15	64.12.138.120 64.12.137.152 64.12.137.121

MX Records:

Preference	TTL	Hostname
15	52m24s	mailin-01.mx.aol.com
15	52m24s	mailin-02.mx.aol.com
15	52m24s	mailin-03.mx.aol.com
15	52m24s	mailin-04.mx.aol.com

Last 5XX Error:

```
-----
```

```
550 REQUESTED ACTION NOT TAKEN: DNS FAILURE
```

```
(at Tue Mar 02 15:17:32 2010 GMT) IP: 10.10.10.10
```

```
-----
```

```

Last TLS Error:                Required - Verify
                                -----
                                TLS required, STARTTLS unavailable
                                (at Tue Mar 02 15:17:32 2010 GMT) IP: 10.10.10.10

Virtual gateway information:
=====

example.com (PublicNet_017):

Host up/down:                  up

Last Activity                   Wed June 22 13:47:02 2005

Recipients                      0

```

**Note**

The Virtual Gateway address information only appears if you are using the `altsrchost` feature.

Determining the Make-up of the Email Queue

To get immediate information about the email queue and determine if a particular recipient host has delivery problems — such as a queue buildup — use the `tophosts` command. The `tophosts` command returns a list of the top 20 recipient hosts in the queue. The list can be sorted by a number of different statistics, including active recipients, connections out, delivered recipients, soft bounced events, and hard bounced recipients. For a description of each item, see [Overview of Managing and Monitoring Using the CLI, page 34-1](#).

Example

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

```
Sort results by:
```

1. Active Recipients
2. Connections Out
3. Delivered Recipients

4. Soft Bounced Events

5. Hard Bounced Recipients

[1]> 1

Status as of: Mon Nov 18 22:22:23 2003

		Active	Conn.	Deliv.	Soft	Hard
#	Recipient Host	Recip	Out	Recip.	Bounced	Bounced
1	aol.com	365	10	255	21	8
2	hotmail.com	290	7	198	28	13
3	yahoo.com	134	6	123	11	19
4	excite.com	98	3	84	9	4
5	msn.com	84	2	76	33	29

mail3.example.com>

Displaying Real-time Activity

The Cisco appliance offers real-time monitoring, which allows you to view the progress of email activity on the system. The `rate` command returns real-time monitoring information about email operations. The information is updated on a periodic interval as specified by you. Use Control-C to stop the `rate` command.

The data shown are listed in [Table 34-7](#)

Table 34-7 Data in the rate Command

Statistic	Description
Connections In	Number of inbound connections.
Connections Out	Number of outbound connections.
Recipients Received	Total number of recipients received into the system.
Recipients Completed	Total number of recipients completed.
Delta	The difference change in Received and Completed recipients since the last data update.
Queue Used	Size of the message queue in kilobytes.

Example

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

Enter the number of seconds between displays.

```
[10]> 1
```

Hit Ctrl-C to return to the main prompt.

Time	Connections		Recipients	Recipients		Queue	
	In	Out	Received	Delta	Completed	Delta	K-Used
23:37:13	10	2	41708833	0	40842686	0	64
23:37:14	8	2	41708841	8	40842692	6	105
23:37:15	9	2	41708848	7	40842700	8	76
23:37:16	7	3	41708852	4	40842705	5	64
23:37:17	5	3	41708858	6	40842711	6	64
23:37:18	9	3	41708871	13	40842722	11	67
23:37:19	7	3	41708881	10	40842734	12	64
23:37:21	11	3	41708893	12	40842744	10	79

^C

The `hostrate` command returns real-time monitoring information about a specific mail host. This information is a subset of the status detail command. (See [Monitoring Detailed Email Status, page 34-8.](#))

Table 34-8 Data in the hostrate Command

Statistic	Description
Host Status	Current status of the specific host: up, down, or unknown.
Current Connections Out	Current number of outbound connections to the host.
Active Recipients in Queue	Total number of active recipients to the specific host in queue.
Active Recipients in Queue Delta	Difference in the total number of active recipients to the specific host in queue since the last known host status.
Delivered Recipients Delta	Difference in the total number of delivered recipients to the specific host in queue since the last known host status.

Table 34-8 Data in the *hostrate* Command

Statistic	Description
Hard Bounced Recipients Delta	Difference in the total number of hard bounced recipients to the specific host in queue since the last known host status.
Soft Bounce Events Delta	Difference in the total number of soft bounced recipients to the specific host in queue since the last known host status.

Use Control-C to stop the *hostrate* command.

Example

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

```
Recipient host:
```

```
[> aol.com
```

```
Enter the number of seconds between displays.
```

```
[10]> 1
```

```

      Time   Host   CrtCncOut  ActvRcp  ActvRcp  DlvRcp  HrdBncRcp  SftBncEvt
          Status                                Delta   Delta   Delta   Delta
23:38:23    up      1          0          0          4          0          0
23:38:24    up      1          0          0          4          0          0
23:38:25    up      1          0          0         12          0          0
^C
```

Monitoring Inbound Email Connections

You may want to monitor hosts that are connecting to the Cisco appliance to identify the large volume senders or to troubleshoot inbound connections to the system. The `topin` command provides a snapshot of the remote hosts connecting to the system. It displays a table with one row for each remote IP address connecting to a specific listener. Two connections from the same IP address to different listeners results in 2 rows in the table. [Table 34-9](#) describes the fields displayed when using the `topin` command.

Table 34-9 Data in the `topin` Command

Statistic	Description
Remote Hostname	Hostname of the remote host, derived from Reverse DNS lookup.
Remote IP Address	IP address of the remote host.
listener	Nickname of the listener on the Cisco appliance that is receiving the connection.
Connections In	The number of concurrent connections from the remote host with the specified IP address open at the time when the command is run.

The system does a reverse DNS lookup to find the remote hostname, and then a forward DNS lookup to validate the name. If the forward lookup does not result in the original IP address, or if the reverse DNS lookup fails, the table displays the IP address in the hostname column. For more information about the process of sender verification, see [Verifying Senders, page 7-28](#).

Example

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

```
Status as of: Sat Aug 23 21:50:54 2003
```

#	Remote hostname	Remote IP addr.	listener	Conn. In
1	mail.remotedomain01.com	172.16.0.2	Incoming01	10
2	mail.remotedomain01.com	172.16.0.2	Incoming02	10
3	mail.remotedomain03.com	172.16.0.4	Incoming01	5
4	mail.remotedomain04.com	172.16.0.5	Incoming02	4
5	mail.remotedomain05.com	172.16.0.6	Incoming01	3
6	mail.remotedomain06.com	172.16.0.7	Incoming02	3
7	mail.remotedomain07.com	172.16.0.8	Incoming01	3

8	mail.remotedomain08.com	172.16.0.9	Incoming01	3
9	mail.remotedomain09.com	172.16.0.10	Incoming01	3
10	mail.remotedomain10.com	172.16.0.11	Incoming01	2
11	mail.remotedomain11.com	172.16.0.12	Incoming01	2
12	mail.remotedomain12.com	172.16.0.13	Incoming02	2
13	mail.remotedomain13.com	172.16.0.14	Incoming01	2
14	mail.remotedomain14.com	172.16.0.15	Incoming01	2
15	mail.remotedomain15.com	172.16.0.16	Incoming01	2
16	mail.remotedomain16.com	172.16.0.17	Incoming01	2
17	mail.remotedomain17.com	172.16.0.18	Incoming01	1
18	mail.remotedomain18.com	172.16.0.19	Incoming02	1
19	mail.remotedomain19.com	172.16.0.20	Incoming01	1
20	mail.remotedomain20.com	172.16.0.21	Incoming01	1

Checking the DNS Status

The `dnsstatus` command returns a counter displaying statistics of DNS lookup and cache information. For each counter, you can view the total number of events since the counter was last reset, since the last system reboot, and over the lifetime of the system.

Table 34-10 lists the available counters.

Table 34-10 Data in the `dnsstatus` Command

Statistic	Description
DNS Requests	A top-level, non-recursive request to the system DNS cache to resolve a domain name.
Network Requests	A request to the network (non-local) to retrieve DNS information.
Cache Hits	A request to the DNS cache where the record was found and returned.
Cache Misses	A request to the DNS cache where the record was not found.

Table 34-10 Data in the *dnsstatus* Command (continued)

Statistic	Description
Cache Exceptions	A request to the DNS cache where the record was found but the domain was unknown.
Cache Expired	A request to the DNS cache where the record was found in the cache, considered for use, and discarded because it was too old. Many entries can exist in the cache even though their time to live (TTL) has been exceeded. As long as these entries are not used, they will not be included in the expires counter. When the cache is flushed, both valid and invalid (too old) entries are deleted. A flush operation does not change the expires counter.

Example

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

```
Status as of: Sat Aug 23 21:57:28 2003
```

Counters:	Reset	Uptime	Lifetime
DNS Requests	211,735,710	8,269,306	252,177,342
Network Requests	182,026,818	6,858,332	206,963,542
Cache Hits	474,675,247	17,934,227	541,605,545
Cache Misses	624,023,089	24,072,819	704,767,877
Cache Exceptions	35,246,211	1,568,005	51,445,744
Cache Expired	418,369	7,800	429,015

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

Resetting Email Monitoring Counters

The `resetcounters` command resets cumulative email monitoring counters. The reset affects global counters as well as per host counters. The reset does not affect the counters on messages in the delivery queue related to retry schedules.



Note

You can also reset the counters in the GUI. See [System Status Page, page 28-30](#).

Example

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

```
Counters reset: Mon Jan 01 12:00:01 2003
```

Identifying Active TCP/IP Services

To identify active TCP/IP services used by your Email Security appliance, use the `tcpservices` command in the command line interface.

Managing the Email Queue

Cisco AsyncOS allows you to perform operations on messages in the email queue. You can delete, bounce, suspend, or redirect messages in the email queue. You can also locate, remove, and archive older messages in your queue.

Deleting Recipients in Queue

If particular recipients are not being delivered or to clear the email queue, use the `deleterecipients` command. The `deleterecipients` command allows you to manage the email delivery queue by deleting specific recipients waiting for delivery. Recipients to be deleted are identified by either the recipient host that the recipient is destined for, or the message sender identified by the specific address given in the Envelope From line of the message envelope. Alternately, you can delete all messages in the delivery queue (all active recipients) at once.

**Note**

To perform the `deleterecipients` function, it is recommended that you place the Cisco appliance in an offline state or suspended delivery (see [Taking an Appliance Offline Using the CLI, page 33-3](#) or [Suspending Email Receiving and Delivery, page 33-2](#)).

**Note**

Although the function is supported in all states, certain messages may be delivered while the function is taking place.

Matches to recipient hosts and senders must be identical string matches. Wild cards are not accepted. The `deleterecipients` command returns the total number of messages deleted. In addition, if a mail log subscription (IronPort text format only) is configured, the message deletion is logged as a separate line.

Example

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

```
Please select how you would like to delete messages:
```

1. By recipient host.
2. By Envelope From address.
3. All.

```
[1]>
```

The Cisco appliance gives you various options to delete recipients depending upon the need. The following example show deleting recipients by recipient host, deleting by Envelope From Address, and deleting all recipients in the queue.

Delete by Recipient Domain

```
Please enter the hostname for the messages you wish to delete.
```

```
[> example.com
```

```
Are you sure you want to delete all messages being delivered to "example.com"? [N]> Y
```

```
Deleting messages, please wait.
```

```
100 messages deleted.
```

Delete by Envelope From Address

```
Please enter the Envelope From address for the messages you wish to delete.

[]> mailadmin@example.com

Are you sure you want to delete all messages with the Envelope From address of
"mailadmin@example.com"? [N]> Y

Deleting messages, please wait.

100 messages deleted.
```

Delete All

```
Are you sure you want to delete all messages in the delivery queue (all active
recipients)? [N]> Y

Deleting messages, please wait.

1000 messages deleted.
```

Bouncing Recipients in Queue

Similar to the `deleterecipients` command, the `bouncerecipients` command allows you to manage the email delivery queue by hard bouncing specific recipients waiting for delivery. Message bouncing follows regular bounce message configuration as specified in the `bounceconfig` command.

**Note**

To perform the `bouncerecipients` function, it is recommended that you place the Cisco appliance in an offline state or suspended delivery (see [Taking an Appliance Offline Using the CLI, page 33-3](#) or [Suspending Email Receiving and Delivery, page 33-2](#)).

**Note**

Although the function is supported in all states, certain messages may be delivered while the function is taking place.

Matches to recipient hosts and senders must be identical string matches. Wild cards are not accepted. The `bouncerecipients` command returns the total number of messages bounced.

**Note**

The `bouncerecipients` function is resource-intensive and may take several minutes to complete. If in offline or suspended delivery state, the actual sending of bounce messages (if hard bounce generation is on) will begin only after Cisco AsyncOS is placed back into the online state by using the `resume` command.

Example

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

```
Please select how you would like to bounce messages:
```

1. By recipient host.
2. By Envelope From address.
3. All.

```
[1]>
```

Recipients to be bounced are identified by either the destination recipient host or the message sender identified by the specific address given in the Envelope From line of the message envelope. Alternately, all messages in the delivery queue can be bounced at once.

Bounce by Recipient Host

```
Please enter the hostname for the messages you wish to bounce.
```

```
[ ]> example.com
```

```
Are you sure you want to bounce all messages being delivered to "example.com"? [N]> Y
```

```
Bouncing messages, please wait.
```

```
100 messages bounced.
```

Bounce by Envelope From Address

```
Please enter the Envelope From address for the messages you wish to bounce.
```

```
[ ]> mailadmin@example.com
```

```
Are you sure you want to bounce all messages with the Envelope From address of  
"mailadmin@example.com"? [N]> Y
```

```
Bouncing messages, please wait.
```

```
100 messages bounced.
```

Bounce All

```
Are you sure you want to bounce all messages in the queue? [N]> y
```

```
Bouncing messages, please wait.
```

```
1000 messages bounced.
```

Redirecting Messages in Queue

The `redirectrecipients` commands allow you to redirect all messages in the email delivery queue to another relay host. Please note that redirecting recipients to a host or IP address that is not prepared to accept large volumes of SMTP mail from this host will cause messages to bounce and possibly result in the loss of mail.



Caution

Redirecting messages to a receiving domain that has `/dev/null` as its destination results in the loss of messages. The CLI does not display a warning if you redirect mail to such a domain. Check the SMTP route for the receiving domain before redirecting messages.

Example

The following example redirects all mail to the `example2.com` host.

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

```
Please enter the hostname or IP address of the machine you want to send all mail to.
```

```
[> example2.com
```

```
WARNING: redirecting recipients to a host or IP address that is not prepared to accept  
large volumes of SMTP mail from this host will cause messages to bounce and possibly  
result in the loss of mail.
```

```
Are you sure you want to redirect all mail in the queue to "example2.com"? [N]> y
```

```
Redirecting messages, please wait.
```

```
246 recipients redirected.
```

Showing Messages Based on Recipient in Queue

Use the `showrecipients` command to show messages from the email delivery queue by recipient host or Envelope From address. You can also show all messages in the queue.

Example

The following example shows messages in the queue for all recipient hosts.

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

```
Please select how you would like to show messages:
```

1. By recipient host.
2. By Envelope From address.
3. All.

```
[1]> 3
```

```
Showing messages, please wait.
```

MID/	Bytes/	Sender/	Subject
[RID]	[Atmps]	Recipient	
1527	1230	user123456@ironport.com	Testing
[0]	[0]	9554@example.com	
1522	1230	user123456@ironport.com	Testing
[0]	[0]	3059@example.com	
1529	1230	user123456@ironport.com	Testing
[0]	[0]	7284@example.com	
1530	1230	user123456@ironport.com	Testing
[0]	[0]	8243@example.com	

```

1532      1230      user123456@ironport.com Testing
[0]      [0]      1820@example.com

1531      1230      user123456@ironport.com Testing
[0]      [0]      9595@example.com

1518      1230      user123456@ironport.com Testing
[0]      [0]      8778@example.com

1535      1230      user123456@ironport.com Testing
[0]      [0]      1703@example.com

1533      1230      user123456@ironport.com Testing
[0]      [0]      3052@example.com

1536      1230      user123456@ironport.com Testing
[0]      [0]      511@example.com

```

Suspending Email Delivery

To temporarily suspend email delivery for maintenance or troubleshooting, use the `suspenddel` command. The `suspenddel` command puts Cisco AsyncOS into suspended delivery state. This state is characterized by the following:

- Outbound email delivery is halted.
- Inbound email connections are accepted.
- Log transfers continue.
- The CLI remains accessible.

The `suspenddel` command lets open outbound connections close, and it stops any new connections from opening. The `suspenddel` command commences immediately, and allows any established connections to successfully close. Use the `resumedel` command to return to regular operations from the suspended delivery state.

**Note**

The “delivery suspend” state is preserved across system reboots. If you use the `suspenddel` command and then reboot the appliance, you must resume delivery after the reboot using the `resumedel` command.

Example

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

Enter the number of seconds to wait before abruptly closing connections.

```
[30]>
```

Waiting for outgoing deliveries to finish...

Mail delivery suspended.

Resuming Email Delivery

The `resumedel` command returns Cisco AsyncOS to normal operating state after using the `suspenddel` command.

Syntax

```
resumedel
```

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

Mail delivery resumed.

Suspending Receiving Email

To temporarily suspend all listeners from receiving email, use the `suspendlistener` command. While receiving is suspended, the system does not accept connections to the specific port of the listener.

This behavior has changed in this release of AsyncOS. In previous releases, the system would accept connections, respond with the following responses and disconnect:

- SMTP: 421 *hostname* Service not available, closing transaction channel
- QMQP: ZService not available

**Note**

The “receiving suspend” state is preserved across system reboots. If you use the `suspendlistener` command and then reboot the appliance, you must use the `resumelister` command before the listener will resume receiving messages.

Syntax

suspendlistener

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

Choose the listener(s) you wish to suspend.

Separate multiple entries with commas.

1. All
2. InboundMail
3. OutboundMail

```
[1]> 1
```

Enter the number of seconds to wait before abruptly closing connections.

```
[30]>
```

Waiting for listeners to exit...

Receiving suspended.

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

Resuming Receiving Email

The `resumelister` command returns Cisco AsyncOS to normal operating state after using the `suspendlistener` command.

Syntax

resumelister

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

Choose the listener(s) you wish to resume.

Separate multiple entries with commas.

1. All
2. InboundMail

```
3. OutboundMail

[1]> 1

Receiving resumed.

mail3.example.com>
```

Resuming Delivery and Receiving of Email

The resume command resumes both delivery and receiving.

Syntax

```
resume

mail3.example.com> resume

Receiving resumed.

Mail delivery resumed.

mail3.example.com>
```

Scheduling Email for Immediate Delivery

Recipients and hosts that are scheduled for later delivery can be immediately retried by using the `delivernow` command. The `delivernow` command allows you to reschedule email in the queue for immediate delivery. All domains that are marked down and any scheduled or soft bounced messages are queued for immediate delivery.

The `delivernow` command can be invoked for all recipients or specific recipients in the queue (scheduled and active). When selecting specific recipients, you must enter the domain name of the recipients to schedule for immediate delivery. The system matches the entire string for character and length.

Syntax

```
delivernow

mail3.example.com> delivernow

Please choose an option for scheduling immediate delivery.

1. By recipient host

2. All messages
```

```
[1]> 1
```

Please enter the domain to schedule for immediate delivery.

```
[ ]> recipient.example.com
```

Rescheduling all messages to recipient.example.com for immediate delivery.

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

Pausing the Work Queue

Processing for LDAP recipient access, masquerading, LDAP re-routing, Message Filters, anti-spam, and the anti-virus scanning engine are all performed in the “work queue.” Refer to [Configuring Routing and Delivery Features, page 24-1](#) for the processing flow and [Table 34-2 on page 34-4](#) for a description of the “Messages in Work Queue” gauge. You can manually pause the work queue portion of message processing using the `workqueue` command.

For example, assume that you wanted to change the configuration of an LDAP server configuration while many messages are in the work queue. Perhaps you want to switch from bouncing to dropping messages based on an LDAP recipient access query. Or perhaps you want to pause the queue while you manually check for the latest anti-virus scanning engine definition files (via the `antivirusupdate` command). The `workqueue` command allows you to pause and resume the work queue to stop processing while you perform other configuration changes.

When you pause and resume the work queue, the event is logged. For example

```
Sun Aug 17 20:01:36 2003 Info: work queue paused, 1900 msgs S
```

```
Sun Aug 17 20:01:39 2003 Info: work queue resumed, 1900 msgs
```

In the following example, the work queue is paused:

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

```
Status as of: Sun Aug 17 20:02:30 2003 GMT
```

```
Status: Operational
```

```
Messages: 1243
```

Choose the operation you want to perform:

```
- STATUS - Display work queue status
```

- PAUSE - Pause the work queue
- RATE - Display work queue statistics over time

[> **pause**

Manually pause work queue? This will only affect unprocessed messages. [N]> **y**

Reason for pausing work queue:

[> **checking LDAP server**

Status as of: Sun Aug 17 20:04:21 2003 GMT

Status: Paused by admin: checking LDAP server

Messages: 1243

**Note**

Entering a reason is optional. If you do not enter a reason, the system logs the reason as “Manually paused by user.”

In this example, the work queue is resumed:

mail3.example.com> **workqueue**

Status as of: Sun Aug 17 20:42:10 2003 GMT

Status: Paused by admin: checking LDAP server

Messages: 1243

Choose the operation you want to perform:

- STATUS - Display work queue status
- RESUME - Resume the work queue
- RATE - Display work queue statistics over time

[> **resume**

```
Status: Operational
```

```
Messages: 1243
```

Locating and Archiving Older Messages

Sometimes older messages remain in the queue because they could not be delivered. You may want to remove and archive these messages. To do this, use the `showmessage` CLI command to display the message for the given message ID. Use the `oldmessage` CLI command to display the oldest non-quarantine message on the system. You can then optionally use the `removemessage` to safely remove the message for the given message ID. This command can only remove messages that are in the work queue, retry queue, or a destination queue. If the message is in none of these queues, it cannot be removed.

You can also use the `archivemessage[mid]` CLI command to archive the message for a given message ID into an mbox file in the configuration directory.

You cannot use the `oldmessage` command to get the message ID for a message in a quarantine. However, if you know the message ID, you can show or archive the specified message. Since the message is not in the work queue, retry queue, or a destination queue, you cannot remove the message with the `removemessage` command.



Note

You cannot perform any of these queue management commands on a message in the Cisco Spam Quarantine.

Syntax

```
archivemessage
```

```
example.com> archivemessage
```

```
Enter the MID to archive and remove.
```

```
[0]> 47
```

```
MID 47 has been saved in file oldmessage_47.mbox in the configuration directory
```

```
example.com>
```

Syntax

```
oldmessage
```

```
example.com> oldmessage
```

```
MID 9: 1 hour 5 mins 35 secs old
```

```
Received: from example.com ([172.16.0.102])
```

```
by example.com with SMTP; 14 Feb 2007 22:11:37 -0800

From: user123@example.com

To: 4031@test.example2.com

Subject: Testing

Message-Id: <20070215061136.68297.16346@example.com>
```

Tracking Messages Within the System

The `findevent` CLI command simplifies the process of tracking messages within the system using the onbox mail log files. The `findevent` CLI command allows you to search through the mail logs for a particular message by searching for a message ID or a regular expression match against the subject header, envelope sender or envelope recipient. You can display results for the current log file, all the log files, or display log files by date. When you view log files by date, you can specify a date or a range of dates.

After you identify the message you want to view logs for, the `findevent` command displays the log information for that message ID including splintering information (split log messages, bounces and system generated messages). The following example shows the `findevent` CLI command tracking the receiving and delivery a message with “confidential” in the subject header:

```
example.com> findevent

Please choose which type of search you want to perform:

1. Search by envelope FROM
2. Search by Message ID
3. Search by Subject
4. Search by envelope TO

[1]> 3

Enter the regular expression to search for.

[]> confidential

Currently configured logs:

1. "mail_logs" Type: "IronPort Text Mail Logs" Retrieval: FTP Poll

Enter the number of the log you wish to use for message tracking.

[]> 1

Please choose which set of logs to search:
```

1. All available log files
2. Select log files by date list
3. Current log file

```
[3]> 3
```

The following matching message IDs were found. Please choose one to show additional log information:

1. MID 4 (Tue Jul 31 17:37:35 2007) sales: confidential

```
[1]> 1
```

```
Tue Jul 31 17:37:32 2007 Info: New SMTP ICID 2 interface Data 1 (172.19.1.86) address
10.251.20.180 reverse dns host unknown verified no
Tue Jul 31 17:37:32 2007 Info: ICID 2 ACCEPT SG None match ALL SBRS None
Tue Jul 31 17:37:35 2007 Info: Start MID 4 ICID 2
Tue Jul 31 17:37:35 2007 Info: MID 4 ICID 2 From: <user@example.com>
Tue Jul 31 17:37:35 2007 Info: MID 4 ICID 2 RID 0 To: <ljohnson@example02.com>
Tue Jul 31 17:37:35 2007 Info: MID 4 Subject 'sales: confidential'
Tue Jul 31 17:37:35 2007 Info: MID 4 ready 4086 bytes from <user@example.com>
Tue Jul 31 17:37:35 2007 Info: MID 4 matched all recipients for per-recipient policy
DEFAULT in the inbound table
Tue Jul 31 17:37:35 2007 Info: ICID 2 close
Tue Jul 31 17:37:37 2007 Info: MID 4 interim verdict using engine: CASE spam negative
Tue Jul 31 17:37:37 2007 Info: MID 4 using engine: CASE spam negative
Tue Jul 31 17:37:37 2007 Info: MID 4 interim AV verdict using Sophos CLEAN
Tue Jul 31 17:37:37 2007 Info: MID 4 antivirus negative
Tue Jul 31 17:37:37 2007 Info: MID 4 queued for delivery
Tue Jul 31 17:37:37 2007 Info: Delivery start DCID 0 MID 4 to RID [0]
Tue Jul 31 17:37:37 2007 Info: Message done DCID 0 MID 4 to RID [0]
Tue Jul 31 17:37:37 2007 Info: MID 4 RID [0] Response '/null'
Tue Jul 31 17:37:37 2007 Info: Message finished MID 4 done
```

SNMP Monitoring

The Cisco AsyncOS operating system supports system status monitoring via SNMP (Simple Network Management Protocol). This includes Cisco'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.) Please note:

- SNMP is **off** by default.
- SNMP SET operations (configuration) are not implemented.
- AsyncOS supports SNMPv1, v2, and v3.
- Message authentication and encryption are mandatory when enabling SNMPv3. Passwords for authentication and encryption should be different. The encryption algorithm can be AES (recommended) or DES. The authentication algorithm can be SHA-1 (recommended) or MD5. The `snmpconfig` command “remembers” your passwords the next time you run the command.
- The SNMPv3 username is: `v3get`.

```
> snmpwalk -v 3 -l AuthNoPriv -u v3get -a MD5 ironport mail.example.com
```

- If you use only SNMPv1 or SNMPv2, you must set a community string. The community string does not default to `public`.
- For SNMPv1 and SNMPv2, you must specify a network from which SNMP GET requests are accepted.
- To use traps, an SNMP manager (not included in AsyncOS) must be running and its IP address entered as the trap target. (You can use a hostname, but if you do, traps will only work if DNS is working.)

Use the `snmpconfig` command to configure SNMP system status for the appliance. After you choose and configure values for an interface, the appliance responds to SNMPv3 GET requests. These version 3 requests must include a matching password. By default, version 1 and 2 requests are rejected. If enabled, version 1 and 2 requests must have a matching community string.

MIB Files

Cisco Systems provides an “enterprise” MIB as well as a “Structure of Management Information” (SMI) file:

- `ASYNOS-MAIL-MIB.txt` — an SNMPv2 compatible description of the Enterprise MIB for Cisco appliances.
- `IRONPORT-SMI.txt` — defines the role of the `ASYNOS-MAIL-MIB` in IronPort’s SNMP managed products.

These files are available on the documentation CD included with your Cisco appliance. You can also request these files through Cisco Customer Support.

Hardware Objects

Hardware sensors conforming to the Intelligent Platform Management Interface Specification (IPMI) report temperature, fan speed, and power supply status.

Table 34-11 shows what hardware derived objects are available for monitoring on what models. The number displayed is the number of instances of that object that can be monitored. For example, you can query the RPMs for 3 fans in the C160/170 appliances and 6 fans in the C300/C600/X1000 appliances.

Table 34-11 **Number of Hardware Objects per Cisco Appliance**

Model	CPU Temp	Ambient Temp	Backplane Temp	Riser Temp	Fans	Power Supply Status	Disk Status	NIC Link
C160/170	1	1	0	0	3	0	2	2
C30/C60	0	0	0	0	0	0	2 (C60 has 4)	3

Table 34-11 Number of Hardware Objects per Cisco Appliance

Model	CPU Temp	Ambient Temp	Backplane Temp	Riser Temp	Fans	Power Supply Status	Disk Status	NIC Link
C300/C600/X1000	2	1	1	1	6	2	4 (C300 has 2)	3 (5 for C600 and X1000 with fiber interface)
C350/C650/X1050	2	1	0	0	4	2	4 (C350 has 2)	3 (5 for the C650 and x1050 with fiber interface)

All models can use SNMP to monitor disk drive health and the link status of Network Interfaces.

Hardware Traps

Table 34-12 lists the temperature and hardware conditions that cause a hardware trap to be sent:

Table 34-12 Hardware Traps: Temperature and Hardware Conditions

Model	High Temp (CPU)	High Temp (Ambient)	High Temp (Backplane)	High Temp (Riser)	Fan Failure	Power Supply	RAID	Link
C160/C170	90C	47C	NA	NA	0 RPMs	Status Change	Status Change	Status Change
C30/C60	NA	NA	NA	NA	NA	NA	Status Change	Status Change
C300/C600/X1000	90C	47C	72C	62C	0 RPMs	Status Change	Status Change	Status Change
C350/C650/X1050	90C	47C	NA	NA	0 RPMs	Status Change	Status Change	Status Change

Status change traps are sent when the status changes. Fan Failure and high temperature traps are sent every 5 seconds. The other traps are failure condition alarm traps — they are sent once when the state changes (healthy to failure). It is a good idea to poll for the hardware status tables and identify possible hardware failures before they become critical. Temperatures within 10 per cent of the critical value may be a cause for concern.

Note that failure condition alarm traps represent a critical failure of the individual component, but may not cause a total system failure. For example, a single fan or power supply can fail on a C600 appliance and the appliance will continue to operate.

SNMP Traps

SNMP provides the ability to send traps, or notifications, to advise an administration application (an SNMP management console, typically) when one or more conditions have been met. Traps are network packets that contain data relating to a component of the system sending the trap. Traps are generated when a condition has been met on the SNMP agent (in this case, the Cisco appliance). After the condition has been met, the SNMP agent then forms an SNMP packet and sends it over port 162, the standard SNMP trap port. In the example below, the trap target of `snmp-monitor.example.com` and the Trap Community string are entered. This is the host running the SNMP management console software that will receive the SNMP traps from the Cisco appliance.

You can configure SNMP traps (enable or disable specific traps) when you enable SNMP for an interface. To specify multiple trap targets: when prompted for the trap target, you may enter up to 10 comma separated IP addresses.

CLI Example

In the following example, the `snmpconfig` command is used to enable SNMP on the “PublicNet” interface on port 161. A passphrase for version 3 is entered and then re-entered for confirmation. The system is configured to service version 1 and 2 requests, and the community string `public` is entered for GET requests from those versions 1 and 2. The trap target of `snmp-monitor.example.com` is entered. Finally, system location and contact information is entered.

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

```
Current SNMP settings:
```

```
SNMP Disabled.
```

```
Choose the operation you want to perform:
```

```
- SETUP - Configure SNMP.
```

```
[ ]> setup
```

```
Do you want to enable SNMP? [N]> y
```

```
Please choose an IP interface for SNMP requests.
```

```
1. Data 1 (192.168.1.1/24: mail3.example.com)
```

```
2. Data 2 (192.168.2.1/24: mail3.example.com)
```

```
3. Management (192.168.44.44/24: mail3.example.com)
```

```
[1]>
```

Enter the SNMPv3 passphrase.

>

Please enter the SNMPv3 passphrase again to confirm.

>

Which port shall the SNMP daemon listen on?

[161]>

Service SNMP V1/V2c requests? [N]> **y**

Enter the SNMP V1/V2c community string.

[> **public**

From which network shall SNMP V1/V2c requests be allowed?

[192.168.2.0/24]>

Enter the Trap target (IP address recommended). Enter "None" to disable traps.

[None]> **10.1.1.29**

Enter the Trap Community string.

[> **tcomm**

Enterprise Trap Status

1. RAIDStatusChange	Enabled
2. fanFailure	Enabled
3. highTemperature	Enabled
4. keyExpiration	Enabled
5. linkDown	Enabled
6. linkUp	Enabled

7. powerSupplyStatusChange Enabled

8. resourceConservationMode Enabled

9. updateFailure Enabled

Do you want to change any of these settings? [N]> **y**

Do you want to disable any of these traps? [Y]>

Enter number or numbers of traps to disable. Separate multiple numbers with commas.

[]> **1,8**

Enterprise Trap Status

1. RAIDStatusChange Disabled

2. fanFailure Enabled

3. highTemperature Enabled

4. keyExpiration Enabled

5. linkDown Enabled

6. linkUp Enabled

7. powerSupplyStatusChange Enabled

8. resourceConservationMode Disabled

9. updateFailure Enabled

Do you want to change any of these settings? [N]>

Enter the System Location string.

[Unknown: Not Yet Configured]> **Network Operations Center - west; rack #31, position 2**

Enter the System Contact string.

[snmp@localhost]> **Joe Administrator, x8888**

Current SNMP settings:

```
Listening on interface "Data 1" 192.168.2.1/24 port 161.
```

```
SNMP v3: Enabled.
```

```
SNMP v1/v2: Enabled, accepting requests from subnet 192.168.2.0/24.
```

```
SNMP v1/v2 Community String: public
```

```
Trap target: 10.1.1.29
```

```
Location: Network Operations Center - west; rack #31, position 2
```

```
System Contact: Joe Administrator, x8888
```

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



SenderBase Network Participation

- [Overview of SenderBase Network Participation, page 35-1](#)
- [Sharing Statistics with SenderBase, page 35-1](#)
- [Frequently Asked Questions, page 35-2](#)

Overview of SenderBase Network Participation

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

Customers participating in the SenderBase Network allow Cisco to collect aggregated email traffic statistics about their organization, increasing the utility of the service for all who use it. Participation is voluntary. Cisco only collects summary data on message attributes and information about how different types of messages were handled by Cisco appliances. For example, Cisco does not collect the message body or the message subject. Personally identifiable information and information that identifies your organization is kept confidential.

Sharing Statistics with SenderBase

Procedure

- Step 1** Go to **Security Services > SenderBase**.
- Step 2** Click **Edit Global Settings**.
- Step 3** Mark the box to enable sharing statistical data with the SenderBase Information Service.
- Checking this box enables the feature globally for the appliance. When enabled, the Context Adaptive Scanning Engine (CASE) is used to collect and report the data (regardless of whether or not Cisco anti-spam scanning is enabled).
- Step 4** (Optional) Enable a proxy server for sharing statistical data with the SenderBase Information Service.
- If you define a proxy server to retrieve rules updates, you can also configure an authenticated username, password, and specific port when connecting to the proxy server in the additional fields provided. To edit these settings, see [System Time, page 33-59](#). You can configure the same settings using the `senderbaseconfig` command in the CLI

Frequently Asked Questions

Cisco recognizes that privacy is important to you, so we design and operate our services with the protection of your privacy in mind. If you enroll in SenderBase Network Participation, Cisco will collect aggregated statistics about your organization's email traffic; however, we do not collect or use any personally identifiable information. Any information Cisco collects that would identify your users or your organization will be treated as confidential.

Why should I participate?

Participating in the SenderBase Network helps us help you. Sharing data with us is important to helping stop email-based threats such as spam, viruses and directory harvest attacks from impacting your organization. Examples of when your participation is especially important include:

- Email attacks that are specifically targeted at your organization, in which case the data you contribute provides the primary source of information to protect you.
- Your organization is one of the first to be hit by a new global email attack, in which case the data you share with us will dramatically improve the speed with which we are able to react to a new threat.

What data do I share?

The data is summarized information on message attributes and information on how different types of messages were handled by Cisco appliances. We do not collect the full body of the message. Again, information provided to Cisco that would identify your users or your organization will be treated as confidential. (See [What does Cisco do to make sure that the data I share is secure?](#), page 35-4 below).

[Table 35-1](#) and [Table 35-2](#) explain a sample log entry in a “human-friendly” format.

Table 35-1 Statistics Shared Per Cisco Appliance

Item	Sample Data
MGA Identifier	MGA 10012
Timestamp	Data from 8 AM to 8:05 AM on July 1, 2005
Software Version Numbers	MGA Version 4.7.0
Rule Set Version Numbers	Anti-Spam Rule Set 102
Anti-virus Update Interval	Updates every 10 minutes
Quarantine Size	500 MB
Quarantine Message Count	50 messages currently in quarantine
Virus Score Threshold	Send messages to quarantine at threat level 3 or higher
Sum of Virus Scores for messages entering quarantine	120
Count of messages entering quarantine	30 (yields average score of 4)
Maximum quarantine time	12 hours
Count of Outbreak quarantine messages broken down by why they entered and exited quarantine, correlated with Anti-Virus result	50 entering quarantine due to .exe rule 30 leaving quarantine due to manual release, and all 30 were virus positive

Table 35-1 Statistics Shared Per Cisco Appliance (continued)

Item	Sample Data
Count of Outbreak quarantine messages broken down by what action was taken upon leaving quarantine	10 messages had attachments stripped after leaving quarantine
Sum of time messages were held in quarantine	20 hours

Table 35-2 Statistics Shared Per IP Address

Item	Sample Data
Message count at various stages within the appliance	Seen by Anti-Virus engine: 100 Seen by Anti-Spam engine: 80
Sum of Anti-Spam and Anti-Virus scores and verdicts	2,000 (sum of anti-spam scores for all messages seen)
Number of messages hitting different Anti-Spam and Anti-Virus rule combinations	100 messages hit rules A and B 50 messages hit rule A only
Number of Connections	20 SMTP Connections
Number of Total and Invalid Recipients	50 total recipients 10 invalid recipients
Hashed Filename(s): (a)	A file <one-way-hash>.pif was found inside an archive attachment called <one-way-hash>.zip.
Obfuscated Filename(s): (b)	A file aaaaaaa0.aaa.pif was found inside a file aaaaaaa.zip.
URL Hostname (c)	There was a link found inside a message to www.domain.com
Obfuscated URL Path (d)	There was a link found inside a message to hostname www.domain.com, and had path aaa000aa/aa00aaa.
Number of Messages by Spam and Virus Scanning Results	10 Spam Positive 10 Spam Negative 5 Spam Suspect 4 Virus Positive 16 Virus Negative 5 Virus Unscannable
Number of messages by different Anti-Spam and Anti-Virus verdicts	500 spam, 300 ham
Count of Messages in Size Ranges	125 in 30K-35K range
Count of different extension types	300 “.exe” attachments

Table 35-2 Statistics Shared Per IP Address

Item	Sample Data (continued)
Correlation of attachment types, true file type, and container type	100 attachments that have a “.doc” extension but are actually “.exe” 50 attachments are “.exe” extensions within a zip
Correlation of extension and true file type with attachment size	30 attachments were “.exe” within the 50-55K range

- (a) Filenames will be encoded in a 1-way hash (MD5).
- (b) Filenames will be sent in an obfuscated form, with all lowercase ASCII letters ([a-z]) replaced with “a,” all uppercase ASCII letters ([A-Z]) replaced with “A,” any multi-byte UTF-8 characters replaced with “x” (to provide privacy for other character sets), all ASCII digits ([0-9]) replaced with “0,” and all other single byte characters (whitespace, punctuation, etc.) maintained. For example, the file Britney1.txt.pif would appear as Aaaaaaa0.aaa.pif.
- (c) URL hostnames point to a web server providing content, much as an IP address does. No confidential information, such as usernames and passwords, are included.
- (d) URL information following the hostname is obfuscated to ensure that any personal information of the user is not revealed.

From AsyncOS 8.5 for Email and later, if IronPort Anti-Spam or Intelligent Multi-Scan feature keys are active and SenderBase Network Participation is enabled, AsyncOS performs the following actions to improve the efficacy of the product:

- Collects information about repetition of certain headers in messages, encrypts the collected information, and adds the encrypted information to the respective messages as headers.
You can submit these processed messages to Cisco for analysis. Each message is reviewed by a team of human analysts and used to enhance the efficacy of the product. For instructions to submit messages to Cisco for analysis, see [Reporting Incorrectly Classified Messages to Cisco Systems, page 13-15](#).
- Sends a random sample of messages to CASE for Antispam scanning, irrespective of their sender's SBRS. CASE scans these messages and uses the results to improve the efficacy of the product. AsyncOS performs this action only when it is idle. As a result, this feedback mechanism does not have any significant impact on the processing of messages.

What does Cisco do to make sure that the data I share is secure?

If you agree to participate in the SenderBase Network:

- Data sent from your Cisco appliances will be sent to the Cisco SenderBase Network servers using the secure protocol HTTPS.
- All customer data will be handled with care at Cisco. This data will be stored in a secure location and access to the data will be limited to employees and contractors at Cisco who require access in order to improve the company's email security products and services or provide customer support.
- No information identifying email recipients or the customer's company will be shared outside of Cisco Systems when reports or statistics are generated based on the data.

Will sharing data impact the performance of my Cisco appliances?

Cisco believes that there will be a minimal performance impact for most customers. We record data that already exists as part of the mail delivery process. Customer data is then aggregated on the appliance and sent to SenderBase servers in batches, typically every 5 minutes. We anticipate that the total size of data transferred via HTTPS will be less than 1% of the bandwidth of a typical company's email traffic.

When enabled, the Context Adaptive Scanning Engine (CASE) is used to collect and report the data (regardless of whether or not Cisco anti-spam scanning is enabled).

**Note**

If you choose to participate in the SenderBase Network, a “body scan” is performed on each message. This happens regardless of whether a filter or other action applied to the message would have triggered a body scan. See [“Body Scanning Rule, page 9-31”](#) for more information about body scanning.

If you have additional questions, please contact Cisco Customer Support. See [Cisco Support Community, page 1-9](#).

Are there other ways I can share data?

For customers wanting to do even more to help Cisco provide top quality security services, there is a command that allows you to share additional data. This higher level of data sharing will also provide attachment filenames in clear, unhashed text, as well as hostnames of URLs in messages. If you are interested in learning more about this feature, please talk to your Systems Engineer or contact Cisco Customer Support.



Other Tasks in the GUI

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

This chapter contains the following sections:

- [The Graphical User Interface \(GUI\), page 36-7](#)
- [System Information in the GUI, page 36-11](#)
- [Gathering XML status from the GUI, page 36-11](#)

The Graphical User Interface (GUI)

After HTTP and/or HTTPS services have been enabled for an interface, you can access the GUI and log in. See the “Accessing the Appliance” chapter for more information.

Enabling the GUI on an Interface

By default, the system ships with HTTP enabled on the Management interface.

To enable the GUI, execute the `interfaceconfig` command at the command-line interface, edit the interface that you want to connect to, and then enable the HTTP services or secure HTTP services, or both.



Note

You can also use the Network > IP Interfaces page to enable or disable the GUI on an interface, once you have the GUI enabled on any other interface. See [IP Interfaces, page A-1](#) for more information.



Note

Enabling secure HTTP on an interface requires you to install a certificate. For more information, see “Enabling a Certificate for HTTPS.”

For either service, you specify the port on which you want the service to be enabled. By default, HTTP is enabled on port 80 and HTTPS on port 443. If you enable both services for an interface, you can automatically redirect HTTP requests to the secure service.

In addition, all users (see [Working with User Accounts, page 32-1](#)) who attempt to access the GUI on this interface (either via HTTP or HTTPS) must authenticate themselves via a standard username and password login page.

**Note**

You must save the changes by using the `commit` command before you are able to access the GUI.

In the following example, the GUI is enabled for the Data 1 interface. The `interfaceconfig` command is used to enable HTTP on port 80 and HTTPS on port 443. (The demonstration certificate is temporarily used for HTTP until the `certconfig` command can be run. For more information, see “Installing Certificates on the Cisco Appliance.”) HTTP requests to port 80 are configured to be automatically redirected to port 443 for the Data1 interface.

Example

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

Currently configured interfaces:

1. Data 1 (192.168.1.1/24 on Data1: mail3.example.com)
2. Data 2 (192.168.2.1/24 on Data2: mail3.example.com)
3. Management (192.168.42.42/24 on Management: mail3.example.com)

Choose the operation you want to perform:

- NEW - Create a new interface.
- EDIT - Modify an interface.
- GROUPS - Define interface groups.
- DELETE - Remove an interface.

```
[> edit
```

Enter the number of the interface you wish to edit.

```
[> 1
```

IP interface name (Ex: "InternalNet"):

```
[Data 1]>
```

Would you like to configure an IPv4 address for this interface (y/n)? [Y]>

IPv4 Address (Ex: 192.168.1.2):

[192.168.1.1]>

Netmask (Ex: "255.255.255.0" or "0xffffffff00"):

[24]>

Would you like to configure an IPv6 address for this interface (y/n)? [N]>

Ethernet interface:

1. Data 1

2. Data 2

3. Management

[1]>

Hostname:

[mail3.example.com]>

Do you want to enable Telnet on this interface? [N]>

Do you want to enable SSH on this interface? [N]>

Do you want to enable FTP on this interface? [N]>

Do you want to enable HTTP on this interface? [N]> **y**

Which port do you want to use for HTTP?

[80]> **80**

Do you want to enable HTTPS on this interface? [N]> **y**

Which port do you want to use for HTTPS?

[443]> **443**

You have not entered a certificate. To assure privacy, run 'certconfig' first. You may use the demo certificate to test HTTPS, but this will not be secure.

Do you really wish to use a demo certificate? [N]> **y**

Both HTTP and HTTPS are enabled for this interface, should HTTP requests redirect to the secure service? [Y]> **y**

Currently configured interfaces:

1. Data 1 (192.168.1.1/24 on Data 1: mail3.example.com)
2. Data 2 (192.168.2.1/24 on Data 2: mail3.example.com)
3. Management (192.168.42.42/24 on Management: mail3.example.com)

Choose the operation you want to perform:

- NEW - Create a new interface.
- EDIT - Modify an interface.
- GROUPS - Define interface groups.
- DELETE - Remove an interface.

[]>

mail3.example.com> **commit**

Please enter some comments describing your changes:

[]> **enabled HTTP, HTTPS for Data 1**

```
Do you want to save the current configuration for rollback? [Y]> n
```

```
Changes committed: Fri May 23 11:42:12 2014 GMT
```

System Information in the GUI

- On the **System Overview** page, you can:
 - View historical graphs and tables showing some of the key system status and performance information.
 - View the version of the AsyncOS operating system installed on the appliance.
 - View a subset of key statistics.
- The **System Status** page provides a detailed representation of all real-time mail and DNS activity for the system. You can also reset the counters for system statistics and view the last time the counters were reset.

Gathering XML status from the GUI

- View status through XML pages, or access XML status information programmatically.

The XML Status feature provides a programmatic method to access email monitoring statistics. Note that some newer browsers can also render XML data directly.

Information from the pages in the GUI in this table is also available as dynamic XML output by accessing the corresponding URL:

GUI Page Name	Corresponding XML status URL
Mail Status	<code>http://hostname/xml/status</code>
Host Mail Status for a Specified Host	<code>http://hostname/xml/hoststatus?hostname=host</code>
DNS Status	<code>http://hostname/xml/dnsstatus</code>
Top Incoming Domains	<code>http://hostname/xml/topin</code>
Top Outgoing Domains^a	<code>http://hostname/xml/tophosts</code>

^a The default sort order for this page is by number of active recipients. You can change the order by appending “?sort=**order**” to the URL, where **order** is **conn_out**, **deliv_recip**, **soft_bounced**, or **hard_bounced**.



Advanced Network Configuration

This chapter includes information about advanced network configuration generally available via the `etherconfig` command, such as NIC pairing, VLANs, Direct Server Return, and more.

- [Media Settings on Ethernet Interfaces, page 37-1](#)
- [Network Interface Card Pairing/Teaming, page 37-3](#)
- [Virtual Local Area Networks \(VLANs\), page 37-6](#)
- [Direct Server Return, page 37-13](#)
- [Ethernet Interface's Maximum Transmission Unit, page 37-18](#)

Media Settings on Ethernet Interfaces

Media settings for the ethernet interfaces can be accessed via the use of the `etherconfig` command. Each ethernet interface is listed with its current setting. By selecting the interface, the possible media settings are displayed. See [Example of Editing Media Settings, page 37-2](#) for an example.

Using etherconfig to Edit Media Settings on Ethernet Interfaces

The `etherconfig` command can be used to set the duplex settings (full/half) as well as the speed (10/100/1000 Mbps) of ethernet interfaces. By default, interfaces automatically select the media settings; however, in some cases you may wish to override this setting.



Note

If you have completed the GUI's System Setup Wizard (or the Command Line Interface `systemsetup` command) as described in the "Setup and Installation" chapter and committed the changes, the default ethernet interface settings should already be configured on your appliance.



Note

Some appliances contain a fiber optic network interface option. If available, you will see two additional ethernet interfaces (Data 3 and Data 4) in the list of available interfaces on these appliances. These gigabit fiber optic interfaces can be paired with the copper (Data 1, Data 2, and Management) interfaces in a heterogeneous configuration. See [Network Interface Card Pairing/Teaming, page 37-3](#).

Example of Editing Media Settings

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

Choose the operation you want to perform:

- MEDIA - View and edit ethernet media settings.
- PAIRING - View and configure NIC Pairing.
- VLAN - View and configure VLANs.
- LOOPBACK - View and configure Loopback.
- MTU - View and configure MTU.

```
[> media
```

Ethernet interfaces:

1. Data 1 (Autoselect: <100baseTX full-duplex>) 00:06:5b:f3:ba:6d
2. Data 2 (Autoselect: <100baseTX full-duplex>) 00:06:5b:f3:ba:6e
3. Management (Autoselect: <100baseTX full-duplex>) 00:02:b3:c7:a2:da

Choose the operation you want to perform:

- EDIT - Edit an ethernet interface.

```
[> edit
```

Enter the name or number of the ethernet interface you wish to edit.

```
[> 2
```

Please choose the Ethernet media options for the Data 2 interface.

1. Autoselect
2. 10baseT/UTP half-duplex
3. 10baseT/UTP full-duplex
4. 100baseTX half-duplex
5. 100baseTX full-duplex

```
6. 1000baseTX half-duplex
```

```
7. 1000baseTX full-duplex
```

```
[1]> 5
```

```
Ethernet interfaces:
```

```
1. Data 1 (Autoselect: <100baseTX full-duplex>) 00:06:5b:f3:ba:6d
```

```
2. Data 2 (100baseTX full-duplex: <100baseTX full-duplex>) 00:06:5b:f3:ba:6e
```

```
3. Management (Autoselect: <100baseTX full-duplex>) 00:02:b3:c7:a2:da
```

```
Choose the operation you want to perform:
```

```
- EDIT - Edit an ethernet interface.
```

```
[ ]>
```

```
Choose the operation you want to perform:
```

```
- MEDIA - View and edit ethernet media settings.
```

```
- PAIRING - View and configure NIC Pairing.
```

```
- VLAN - View and configure VLANs.
```

```
- LOOPBACK - View and configure Loopback.
```

```
- MTU - View and configure MTU.
```

```
[ ]>
```

Network Interface Card Pairing/Teaming

NIC pairing allows you to combine any two physical data ports in order to provide a backup Ethernet interface if the data path from the NIC to the upstream Ethernet port should fail. Basically, pairing configures the Ethernet interfaces so that there is a primary interface and a backup interface. If the primary interface fails (i.e. if the carrier between the NIC and the upstream node is disrupted), the backup interface becomes active and an alert is sent. When the primary interface is up again, this interface will become active automatically. Within the documentation for this product, NIC pairing is synonymous with NIC teaming.

**Note**

NIC pairing is not available on Email Security virtual appliances.

You can create more than one NIC pair, providing you have enough data ports. When creating pairs, you can combine any two data ports. For example:

- Data 1 and Data 2
- Data 3 and Data 4
- Data 2 and Data 3
- etc.

Some Cisco appliances contain a fiber optic network interface option. If available, you will see two additional ethernet interfaces (Data 3 and Data 4) in the list of available interfaces on these appliances. These gigabit fiber optic interfaces can be paired with the copper (Data 1, Data 2, and Management) interfaces in a heterogeneous configuration.

NIC Pairing and VLANs

VLANs (see [Virtual Local Area Networks \(VLANs\)](#), page 37-6) are only allowed on the primary interface.

NIC Pair Naming

When creating NIC pairs, you must specify a name to use to refer to the pair. NIC pairs created in versions of AsyncOS prior to version 4.5 will automatically receive the default name of 'Pair 1' following an upgrade.

Any alerts generated regarding NIC pairing will reference the specific NIC pair by name.

NIC Pairing and Existing Listeners

If you enable NIC pairing on an interface that has listeners assigned to it, you are prompted to either delete, reassign, or disable all listeners assigned to the backup interface.

Enabling NIC Pairing via the etherconfig Command



Note

NIC pairing is not available on Email Security virtual appliances.

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

Choose the operation you want to perform:

- MEDIA - View and edit ethernet media settings.
- PAIRING - View and configure NIC Pairing.
- VLAN - View and configure VLANs.
- LOOPBACK - View and configure Loopback.

- MTU - View and configure MTU.

[> **pairing**

Paired interfaces:

Choose the operation you want to perform:

- NEW - Create a new pairing.

[> **new**

Please enter a name for this pair (Ex: "Pair 1"):

[> **Pair 1**

Warning: The backup (Data 2) for the NIC Pair is currently configured with one or more IP addresses. If you continue, the Data 2 interface will be deleted.

Do you want to continue? [N]> y

The interface you are deleting is currently used by listener "OutgoingMail".

What would you like to do?

1. Delete: Remove the listener and all its settings.
2. Change: Choose a new interface.
3. Ignore: Leave the listener configured for interface "Data 2" (the listener will be disabled until you add a new interface named "Data 2" or edit the listener's settings).

[1]>

Listener OutgoingMail deleted for mail3.example.com.

Interface Data 2 deleted.

Paired interfaces:

1. Pair 1:

Primary (Data 1) Active, Link is up

Backup (Data 2) Standby, Link is up


Choose the operation you want to perform:

- DELETE - Delete a pairing.
- STATUS - Refresh status.

[]>

Virtual Local Area Networks (VLANs)

VLANs are virtual local area networks bound to physical data ports. You can configure VLANs to increase the number of networks the appliance can connect to beyond the number of physical interfaces included. For example, some appliance models have three interfaces: Data 1, Data 2, and Management. VLANs allow more networks to be defined on separate “ports” on existing listeners. (See [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) for more information.) You can configure multiple VLANs on any physical network port. [Figure 37-1](#) provides an example of configuring several VLANs on the Data 2 interface.

Figure 37-1 Using VLANs to increase the number of networks available on the appliance

VLANs can be used to segment networks for security purposes, to ease administration, or increase bandwidth. VLANs appear as dynamic “Data Ports” labeled in the format of: “VLAN DDDD” where the “DDDD” is the ID and is an integer up to 4 digits long (VLAN 2, or VLAN 4094 for example). AsyncOS supports up to 30 VLANs. Duplicate VLAN IDs are not allowed on your appliance.


VLANs and Physical Ports

A physical port does not need an IP address configured in order to be in a VLAN. The physical port on which a VLAN is created can have an IP that will receive non-VLAN traffic, so you can have both VLAN and non-VLAN traffic on the same interface.

VLANs can be created on all “Data” and “Management” ports, including fiber optic data ports available on some appliance models.

VLANs can be used with NIC pairing (available on paired NICs) and with Direct Server Return (DSR).

[Figure 37-2](#) illustrates a use case showing how two mail servers unable to communicate directly due to VLAN limitations can send mail through the Email Security appliance. The blue line shows mail coming from the sales network (VLAN1) to the appliance. The appliance will process the mail as normal and then, upon delivery, tag the packets with the destination VLAN information (red line).

Figure 37-2 Using VLANs to Facilitate Communication Between Appliances

Managing VLANs

You can create, edit and delete VLANs via the `etherconfig` command. Once created, a VLAN can be configured via the Network -> Interfaces page or the `interfaceconfig` command in the CLI. Remember to commit all changes.

Creating a New VLAN via the etherconfig Command

In this example, two VLANs are created (named VLAN 31 and VLAN 34) on the Data 1 port:

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

Choose the operation you want to perform:

- MEDIA - View and edit ethernet media settings.
- PAIRING - View and configure NIC Pairing.
- VLAN - View and configure VLANs.
- LOOPBACK - View and configure Loopback.
- MTU - View and configure MTU.

```
[> vlan
```

VLAN interfaces:

Choose the operation you want to perform:

- NEW - Create a new VLAN.

[> **new**

VLAN ID for the interface (Ex: "34"):

[> **34**

Enter the name or number of the ethernet interface you wish bind to:

1. Data 1
2. Data 2
3. Management

[1]> **1**

VLAN interfaces:

1. VLAN 34 (Data 1)

Choose the operation you want to perform:

- NEW - Create a new VLAN.
- EDIT - Edit a VLAN.
- DELETE - Delete a VLAN.

[> **new**

VLAN ID for the interface (Ex: "34"):

[> **31**

Enter the name or number of the ethernet interface you wish bind to:

1. Data 1
2. Data 2
3. Management

```
[1]> 1
```

VLAN interfaces:

1. VLAN 31 (Data 1)
2. VLAN 34 (Data 1)

Choose the operation you want to perform:

- NEW - Create a new VLAN.
- EDIT - Edit a VLAN.
- DELETE - Delete a VLAN.

```
[>
```

Choose the operation you want to perform:

- MEDIA - View and edit ethernet media settings.
- PAIRING - View and configure NIC Pairing.
- VLAN - View and configure VLANs.
- LOOPBACK - View and configure Loopback.
- MTU - View and configure MTU.

```
[>
```

Creating an IP Interface on a VLAN via the interfaceconfig Command

In this example, a new IP interface is created on the VLAN 31 ethernet interface.



Note

Making changes to an interface may close your connection to the appliance.

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

Currently configured interfaces:

1. Data 1 (10.10.1.10/24: example.com)
2. Management (10.10.0.10/24: example.com)

Choose the operation you want to perform:

- NEW - Create a new interface.
- EDIT - Modify an interface.
- GROUPS - Define interface groups.
- DELETE - Remove an interface.

[> **new**

Please enter a name for this IP interface (Ex: "InternalNet"):

[> **InternalVLAN31**

Would you like to configure an IPv4 address for this interface (y/n)? [Y]>

IPv4 Address (Ex: 10.10.10.10):

[> **10.10.31.10**

Netmask (Ex: "255.255.255.0" or "0xffffffff"):

[255.255.255.0]>

Would you like to configure an IPv6 address for this interface (y/n)? [N]>

Ethernet interface:

1. Data 1
2. Data 2
3. Management
4. VLAN 31
5. VLAN 34

[1]> **4**

Hostname:

```
[> mail31.example.com
```

Do you want to enable Telnet on this interface? [N]>

Do you want to enable SSH on this interface? [N]>

Do you want to enable FTP on this interface? [N]>

Do you want to enable HTTP on this interface? [N]>

Do you want to enable HTTPS on this interface? [N]>

Currently configured interfaces:

1. Data 1 (10.10.1.10/24: example.com)
2. InternalVLAN31 (10.10.31.10/24: mail31.example.com)
3. Management (10.10.0.10/24: example.com)

Choose the operation you want to perform:

- NEW - Create a new interface.
- EDIT - Modify an interface.
- GROUPS - Define interface groups.
- DELETE - Remove an interface.

```
[>
```

You can also configure VLANs via the Network -> Listeners page:

Figure 37-3 Using a VLAN when Creating a New IP Interface via the GUI
Add IP Interface

IP Interface Settings													
Name:	Internal/VLAN31												
Ethernet Port:	VLAN 31												
IP Address:	10.10.31.10												
Netmask:	255.255.255.0												
Hostname:	mail31.example.com												
Services:	<table border="1"> <thead> <tr> <th>Service</th> <th>Port</th> </tr> </thead> <tbody> <tr> <td>☐ FTP</td> <td>21</td> </tr> <tr> <td>☐ Telnet</td> <td>23</td> </tr> <tr> <td>☐ SSH</td> <td>22</td> </tr> <tr> <td>☐ HTTP</td> <td>80</td> </tr> <tr> <td>☐ HTTPS</td> <td>443</td> </tr> </tbody> </table>	Service	Port	☐ FTP	21	☐ Telnet	23	☐ SSH	22	☐ HTTP	80	☐ HTTPS	443
Service	Port												
☐ FTP	21												
☐ Telnet	23												
☐ SSH	22												
☐ HTTP	80												
☐ HTTPS	443												
Redirect HTTP Requests to HTTPS:	☐ Enable Redirect (HTTP and HTTPS Services will be turned on)												
Cancel Submit													

Direct Server Return

Direct Server Return (DSR) is a way of providing support for a light-weight load balancing mechanism to load balance between multiple Email Security appliances sharing the same Virtual IP (VIP).

DSR is implemented via an IP interface created on the “loopback” ethernet interface on the appliance.



Note

Configuring load balancing for Email Security appliances is beyond the scope of this document.

Enabling Direct Server Return

Enable DSR by enabling the “loopback” ethernet interface on each participating appliance. Next, create an IP interface on the loopback interface with a virtual IP (VIP) via the `interfaceconfig` command in the CLI or via the Network -> Interfaces page in the GUI. Finally, create a listener on the new IP interface via the `listenerconfig` command in the CLI or via the Network -> Listeners page in the GUI. Remember to commit all changes.




Note

Using the loopback interface prevents the appliance from issuing ARP replies for that specific interface.

When enabling DSR, the following rules apply:

- All systems use the same Virtual IP (VIP) address
- All systems must be on the same switch and subnet as the load balancer

Figure 37-4 Using DSR to Load Balance Between Multiple Email Security Appliances on a Switch

Enabling the Loopback Interface via the etherconfig Command

Once enabled, the loopback interface is treated like any other interface (e.g. Data 1):

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

Choose the operation you want to perform:

- MEDIA - View and edit ethernet media settings.
- PAIRING - View and configure NIC Pairing.
- VLAN - View and configure VLANs.
- LOOPBACK - View and configure Loopback.
- MTU - View and configure MTU.

```
[>] loopback
```

Currently configured loopback interface:

Choose the operation you want to perform:

- ENABLE - Enable Loopback Interface.

```
[>] enable
```

Currently configured loopback interface:

1. Loopback

Choose the operation you want to perform:

- DISABLE - Disable Loopback Interface.

[]>

Choose the operation you want to perform:

- MEDIA - View and edit ethernet media settings.

- PAIRING - View and configure NIC Pairing.

- VLAN - View and configure VLANs.

- LOOPBACK - View and configure Loopback.

- MTU - View and configure MTU.

[]>

Creating an IP Interface on Loopback via the interfaceconfig Command

Create an IP interface on the loopback interface:

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

Currently configured interfaces:

1. Data 1 (10.10.1.10/24: example.com)

2. InternalV1 (10.10.31.10/24: mail31.example.com)

3. Management (10.10.0.10/24: example.com)

Choose the operation you want to perform:

- NEW - Create a new interface.

- EDIT - Modify an interface.

- GROUPS - Define interface groups.

- DELETE - Remove an interface.


```
[> new
```

Please enter a name for this IP interface (Ex: "InternalNet"):

```
[> LoopVIP
```

Would you like to configure an IPv4 address for this interface (y/n)? [Y]>

IPv4 Address (Ex: 10.10.10.10):

```
[> 10.10.1.11
```

Netmask (Ex: "255.255.255.0" or "0xffffffff00"):

```
[255.255.255.0]> 255.255.255.255
```

Would you like to configure an IPv6 address for this interface (y/n)? [N]>

Ethernet interface:

1. Data 1
2. Data 2
3. Loopback
4. Management
5. VLAN 31
6. VLAN 34

```
[1]> 3
```

Hostname:

```
[> example.com
```

Do you want to enable Telnet on this interface? [N]>

```
Do you want to enable SSH on this interface? [N]>

Do you want to enable FTP on this interface? [N]>

Do you want to enable HTTP on this interface? [N]>

Do you want to enable HTTPS on this interface? [N]>

Currently configured interfaces:

1. Data 1 (10.10.1.10/24: example.com)
2. InternalV1 (10.10.31.10/24: mail31.example.com)
3. LoopVIP (10.10.1.11/24: example.com)
4. Management (10.10.0.10/24: example.com)

Choose the operation you want to perform:

- NEW - Create a new interface.
- EDIT - Modify an interface.
- GROUPS - Define interface groups.
- DELETE - Remove an interface.

[ ]>
```

Creating a Listener on the New IP Interface

Create a listener on the new IP interface via the GUI or the CLI. For example, [Figure 37-5](#) shows the newly created IP interface available in the Add Listener page in the GUI.

Figure 37-5 Creating a Listener on the New Loopback IP Interface
Add Listener

Listener Settings	
Name:	
Type of Listener:	☒ Public ☐ Private
Interface:	Data 1 (10.10.1.10/24: example.com) ▼
Bounce Profile:	Data 1 (10.10.1.10/24: example.com) InternalV1 (10.10.31.10/24: mail31.example.com) LoopVIP (10.10.11.10/24: mail11.example.com) Management (10.10.2.10/24: example.com)
Disclaimer Above:	Disclaimer text will be applied above the message body.
Disclaimer Below:	None ▼ Disclaimer text will be applied below the message body.
SMTP Authentication Profile:	None ▼
Certificate:	System Default ▼

Ethernet Interface's Maximum Transmission Unit

The maximum transmission unit (MTU) is the largest unit of data that an ethernet interface will accept. You can decrease the MTU for an ethernet interface using the `etherconfig` command. The default MTU size is 1500 bytes, which is the largest MTU that the ethernet interface can accept.

To edit an interface's MTU:

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

Choose the operation you want to perform:

- MEDIA - View and edit ethernet media settings.
- PAIRING - View and configure NIC Pairing.
- VLAN - View and configure VLANs.
- LOOPBACK - View and configure Loopback.
- MTU - View and configure MTU.

```
[> mtu
```

Ethernet interfaces:

1. Data 1 mtu 1400
2. Data 2 default mtu 1500
3. Management default mtu 1500

Choose the operation you want to perform:

- EDIT - Edit an ethernet interface.

[> **edit**

Enter the name or number of the ethernet interface you wish to edit.

[> 2

Please enter a non-default (1500) MTU value for the Data 2 interface.

[> **1200**

Ethernet interfaces:

1. Data 1 mtu 1400
2. Data 2 mtu 1200
3. Management default mtu 1500

Choose the operation you want to perform:

- EDIT - Edit an ethernet interface.

[>



Logging

- [Overview, page 38-1](#)
- [Log Types, page 38-7](#)
- [Log Subscriptions, page 38-38](#)

Overview

- [Understanding Log Files and Log Subscriptions, page 38-1](#)
- [Log Types, page 38-1](#)
- [Log Retrieval Methods, page 38-6](#)

Understanding Log Files and Log Subscriptions

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

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

Cisco offers the M-Series Content Security Management appliance for centralized reporting and tracking tool for logs from multiple Email Security appliances. See your Cisco representative for more information.

A log subscription associates a log type with a name, logging level, and other constraints such as size and destination information; multiple subscriptions for the same log type are permitted.

Log Types

The log type indicates what information will be recorded within the generated log such as message data, system statistics, binary or textual data. You select the log type when creating a log subscription. See [Log Subscriptions, page 38-38](#) for more information.

AsyncOS for Email generates the following log types:

Table 38-1 Log Types

Log	Description
Text Mail Logs	Text mail logs record information regarding the operations of the email system. For example, message receiving, message delivery attempts, open and closed connections, bounces, TLS connections, and others.
qmail Format Mail Logs	qmail format delivery logs record the same information regarding the operations of the email system as delivery logs following, but stored in qmail format.
Delivery Logs	Delivery logs record critical information about the email delivery operations of the Email Security appliance — for example, information regarding each recipient delivery and bounce at the time of the delivery attempt. The log messages are “stateless,” meaning that all associated information is recorded in each log message and users need not reference previous log messages for information about the current delivery attempt. Delivery logs are recorded in a binary format for resource efficiency. Delivery Log files must be post-processed using a provided utility to convert them to XML or CSV (comma-separated values) format. The conversion tools are located at: http://support.ironport.com
Bounce Logs	Bounce logs record information about bounced recipients. The information recorded for each bounced recipient includes: the message ID, the recipient ID, the Envelope From address, the Envelope To address, the reason for the recipient bounce, and the response code from the recipient host. In addition, you can choose to log a fixed amount of each bounced recipient message. This amount is defined in bytes and the default is zero.
Status Logs	This log file records system statistics found in the CLI status commands, including <code>status detail</code> and <code>dnsstatus</code> . The period of recording is set using the <code>setup</code> subcommand in <code>logconfig</code> . Each counter or rate reported in status logs is the value since the last time the counter was reset.
Domain Debug Logs	Domain debug logs record the client and server communication during an SMTP conversation between the Email Security appliance and a specified recipient host. This log type can be used to debug issues with specific recipient hosts. You must specify the total number of SMTP sessions to record in the log file. As sessions are recorded, this number decreases. You can stop domain debug before all sessions have been recorded by deleting or editing the log subscription.
Injection Debug Logs	Injection debug logs record the SMTP conversation between the Email Security appliance and a specified host connecting to the system. Injection debug logs are useful for troubleshooting communication problems between the Email Security appliance and a host on the Internet.
System Logs	System logs record the following: boot information, virtual appliance license expiration alerts, DNS status information, and comments users typed using <code>commit</code> command. System logs are useful for troubleshooting the basic state of the appliance.
CLI Audit Logs	The CLI audit logs record all CLI activity on the system.
FTP Server Logs	FTP logs record information about the FTP services enabled on the interface. Connection details and user activity are recorded.

Table 38-1 Log Types (continued)

Log	Description
GUI Logs	See HTTP Logs.
HTTP Logs	HTTP logs record information about the HTTP and/or secure HTTP services enabled on the interface. Because the graphical user interface (GUI) is accessed via HTTP, the HTTP logs are ostensibly the GUI equivalent of the CLI Audit logs. Session data (new session, session expired) and pages accessed in the GUI are recorded. These logs also include information about SMTP transactions, for example information about scheduled reports emailed from the appliance.
NTP Logs	NTP logs record the conversation between the appliance and any NTP (Network Time Protocol) servers configured. For more information, see “Editing the Network Time Protocol (NTP) Configuration (Time Keeping Method)” in the “System Administration” chapter.
LDAP Debug Logs	LDAP debug logs are meant for debugging LDAP installations. (See the “LDAP Queries” chapter.) Useful information about the queries that the Email Security appliance is sending to the LDAP server are recorded here.
Anti-Spam Logs	Anti-spam logs record the status of the anti-spam scanning feature of your system, including the status on receiving updates of the latest anti-spam rules. Also, any logs related to the Context Adaptive Scanning Engine are logged here.
Anti-Spam Archive	If you enabled an Anti-Spam scanning feature, messages that are scanned and associated with the “archive message” action are archived here. The format is an mbox-format log file. For more information about anti-spam engines, see the “Anti-Spam” chapter.
Anti-Virus Logs	AntiVirus logs record the status of the anti-virus scanning feature of your system, including the status on receiving updates of the latest anti-virus identity files.
Anti-Virus Archive	If you enabled an anti-virus engine, messages that are scanned and associated with the “archive message” action are archived here. The format is an mbox-format log file. For more information, see the “Anti-Virus” chapter.
AMP Engine Logs	Logging about the status of Advanced Malware Protection features.
AMP Archive	If you have configured mail policies to archive messages that Advanced Malware Protection has found to have attachments that are unscannable or contain malware, those messages are archived here. The format is an mbox-format log file.
Scanning Logs	The scanning log contains all LOG and COMMON messages for scanning engines (see Alerts, page 33-34). This is typically application faults, alert sent, alert failed, and log error messages. This log does not apply to system-wide alerts.
Spam Quarantine Logs	Spam Quarantine logs record actions associated with the Spam Quarantine processes.
Spam Quarantine GUI Logs	Spam Quarantine logs record actions associated with the Spam Quarantine including configuration via the GUI, end user authentication, and end user actions (releasing email, etc.).

Table 38-1 **Log Types (continued)**

Log	Description
SMTP Conversation Logs	The SMTP conversation log records all parts of incoming and outgoing SMTP conversations.
Safe/Block Lists Logs	Safelist/blocklist logs record data about the safelist/blocklist settings and database.
Reporting Logs	Reporting logs record actions associated with the processes of the centralized reporting service.
Reporting Query Logs	Reporting query logs record actions associated with the reporting queries that are run on the appliance.
Updater Logs	The updater log records events related to updates for system services, such as McAfee Anti-Virus definition updates.
Tracking Logs	Tracking logs record actions associated with the processes of the tracking service. Tracking logs are a subset of the mail logs.
Authentication Logs	The authentication log records successful user logins and unsuccessful login attempts.
Configuration History Logs	Configuration history logs record the following information: What changes were made on the Email Security appliance, and when were the changes made? A new configuration history log is created each time a user commits a change.
Upgrade Logs	Status information about upgrade download and installation.
API Logs	API logs record various events related to Cisco AsyncOS API for Email, for example: • API has started or stopped • Connection to the API failed or closed (after providing response) • Authentication succeeded or failed • Request contains errors • Error while communicating network configuration changes with AsyncOS API

Log Type Characteristics

Table 38-2 summarizes the different characteristics of each log type.

Table 38-2 Log Type Comparison

	Contains												
	Transactional	Stateless	Recorded as text	Recorded as mbox file	Recorded as binary	Periodic Status Information	Message Receiving Information	Delivery Information	Individual Hard Bounces	Individual Soft Bounces	Injection SMTP Conversation	Header Logging	Delivery SMTP Conversation
Mail Logs	•		•			•	•	•	•	•		•	
qmail Format Delivery Logs		•			•		•	•	•			•	
Delivery Log		•			•		•	•	•			•	
Bounce Logs	•		•						•	•		•	
Status Logs		•	•			•							
Domain Debug Logs	•		•					•	•	•			•
Injection Debug Logs	•		•				•				•		
System Logs	•		•			•							
CLI Audit Logs	•		•			•							
FTP Server Logs	•		•			•							
HTTP Logs	•		•			•							
NTP Logs	•		•			•							
LDAP Logs	•		•										
Anti-spam logs	•		•			•							
Anti-Spam Archive Logs				•									
Anti-virus Logs	•		•			•							
Anti-Virus Archive				•									
Scanning Logs	•		•			•							•
Spam Quarantine	•		•			•							
Spam Quarantine GUI	•		•			•							

Table 38-2 Log Type Comparison (continued)

	Contains												
	Transactional	Stateless	Recorded as text	Recorded as mbox file	Recorded as binary	Periodic Status Information	Message Receiving Information	Delivery Information	Individual Hard Bounces	Individual Soft Bounces	Injection SMTP Conversation	Header Logging	Delivery SMTP Conversation
Safe/Block Lists Logs	•		•			•							
Reporting Logs	•		•		•								
Reporting Query Logs	•		•		•								
Updater Logs			•										
Tracking Logs	•				•	•	•	•	•	•		•	
Authentication Logs	•		•										
Configuration History Logs	•		•										•
API Logs	•		•										

Log Retrieval Methods

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

Table 38-3 Log Transfer Protocols

Manually Download	This method lets you access log files at any time by clicking a link to the log directory on the Log Subscriptions page, then clicking the log file to access. Depending on your browser, you can view the file in a browser window, or open or save it as a text file. This method uses the HTTP(S) protocol and is the default retrieval method. Note Using this method, you cannot retrieve logs for any computer in a cluster, regardless of level (machine, group, or cluster), even if you specify this method in the CLI.
FTP Push	This method periodically pushes log files to an FTP server on a remote computer. The subscription requires a username, password, and destination directory on the remote computer. Log files are transferred based on a rollover schedule set by you.

Table 38-3 Log Transfer Protocols (continued)

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

Log Filenames and Directory Structure

AsyncOS creates a directory for each log subscription based on the log subscription name. The actual name of the log file in the directory is composed of the log filename specified by you, the timestamp when the log file was started, and a single-character status code. The filename of logs are made using the following formula:

```
/LogSubscriptionName/LogFilename.@timestamp.statuscode
```

Status codes may be `.current` or `.s` (signifying saved). You should only transfer or delete log files with the saved status.

Log Rollover and Transfer Schedule

Log files are created by log subscriptions, and are rolled over (and transferred, if a push-based retrieval option is selected) based on the first user-specified condition reached: maximum file size or scheduled rollover. Use the `logconfig` command in the CLI or the Log Subscriptions page in the GUI to configure both the maximum file size and time interval for scheduled rollovers. You can also use the **Rollover Now** button in the GUI or the `rollovernow` command in the CLI to rollover selected log subscriptions. See [Rolling Over Log Subscriptions, page 38-43](#) for more information on scheduling rollovers.

Logs retrieved using manual download are saved until they reach the maximum number you specify (the default is 10 files) or until the system needs more space for log files.

Logs Enabled by Default

Your Email Security appliance is pre-configured with many log subscriptions enabled by default (other logs may be configured depending on which license keys you have applied). By default, the retrieval method is “Manually Download.”

All pre-configured log subscriptions have a Log Level of 3, except for `error_logs` which is set at 1 so that it will contain only errors. See [Log Levels, page 38-39](#) for more information. For information about creating new log subscriptions, or modifying existing ones, see [Log Subscriptions, page 38-38](#).

Log Types

- [Using Text Mail Logs, page 38-8](#)
- [Using Delivery Logs, page 38-15](#)

- [Using Bounce Logs, page 38-17](#)
- [Using Status Logs, page 38-19](#)
- [Using Domain Debug Logs, page 38-22](#)
- [Using Injection Debug Logs, page 38-23](#)
- [Using System Logs, page 38-24](#)
- [Using CLI Audit Logs, page 38-25](#)
- [Using FTP Server Logs, page 38-26](#)
- [Using HTTP Logs, page 38-27](#)
- [Using NTP Logs, page 38-28](#)
- [Using Scanning Logs, page 38-28](#)
- [Using Anti-Spam Logs, page 38-29](#)
- [Using Anti-Virus Logs, page 38-29](#)
- [Using Spam Quarantine Logs, page 38-30](#)
- [Using Spam Quarantine GUI Logs, page 38-30](#)
- [Using LDAP Debug Logs, page 38-31](#)
- [Using Safelist/Blocklist Logs, page 38-32](#)
- [Using Reporting Logs, page 38-33](#)
- [Using Reporting Query Logs, page 38-34](#)
- [Using Updater Logs, page 38-35](#)
- [Understanding Tracking Logs, page 38-36](#)
- [Using Authentication Logs, page 38-37](#)
- [Using Configuration History Logs, page 38-37](#)

Timestamps in Log Files

The following log files include the begin and end date of the log itself, the version of AsyncOS, and the GMT offset (provided in seconds, and only at the beginning of the log):

- Anti-Virus log
- LDAP log
- System log
- Mail log

Using Text Mail Logs

They contain details of email receiving, email delivery and bounces. Status information is also written to the mail log every minute. These logs are a useful source of information to understand delivery of specific messages and to analyze system performance.

These logs do not require any special configuration. However, you must configure the system properly to view attachment names, and attachment names may not always be logged. For information, see [Enabling Message Tracking, page 29-1](#) and [Message Tracking Overview, page 29-1](#).

Information displayed in text mail logs is shown in [Table 38-4](#).

Table 38-4 *Text Mail Log Statistics*

Statistic	Description
ICID	Injection Connection ID. This is a numerical identifier for an individual SMTP connection to the system, over which 1 to thousands of individual messages may be sent.
DCID	Delivery Connection ID. This is a numerical identifier for an individual SMTP connection to another server, for delivery of 1 to thousands of messages, each with some or all of their RIDs being delivered in a single message transmission.
RCID	RPC Connection ID. This is a numerical identifier for an individual RPC connection to the Spam quarantine. It is used to track messages as they are sent to and from the Spam Quarantine.
MID	Message ID: Use this to track messages as they flow through the logs.
RID	Recipient ID: Each message recipient is assigned an ID.
New	New connection initiated.
Start	New message started.

Interpreting a Text Mail Log

Use the following sample as a guide to interpret log files.



Note

Individual lines in log files are NOT numbered. They are numbered here only for sample purposes.

Table 38-5 *Text Mail Log Detail*

1	Mon Apr 17 19:56:22 2003 Info: New SMTP ICID 5 interface Management (10.1.1.1) address 10.1.1.209 reverse dns host remotehost.com verified yes
2	Mon Apr 17 19:57:20 2003 Info: Start MID 6 ICID 5
3	Mon Apr 17 19:57:20 2003 Info: MID 6 ICID 5 From: <sender@remotehost.com>
4	Mon Apr 17 19:58:06 2003 Info: MID 6 ICID 5 RID 0 To: <mary@yourdomain.com>
5	Mon Apr 17 19:59:52 2003 Info: MID 6 ready 100 bytes from <sender@remotehost.com>
6	Mon Apr 17 19:59:59 2003 Info: ICID 5 close
7	Mon Mar 31 20:10:58 2003 Info: New SMTP DCID 8 interface 192.168.42.42 address 10.5.3.25
8	Mon Mar 31 20:10:58 2003 Info: Delivery start DCID 8 MID 6 to RID [0]
9	Mon Mar 31 20:10:58 2003 Info: Message done DCID 8 MID 6 to RID [0]
10	Mon Mar 31 20:11:03 2003 Info: DCID 8 close

Use [Table 38-6](#) as a guide to reading the preceding log file.

Table 38-6 *Detail of Text Mail Log Example*

Line Number	Description
1.	A new connection is initiated into the system and assigned an Injection ID (ICID) of “5.” The connection was received on the Management IP interface and was initiated from the remote host at 10.1.1.209.
2.	The message was assigned a Message ID (MID) of “6” after the MAIL FROM command is issued from the client.
3.	The sender address is identified and accepted.
4.	The recipient is identified and assigned a Recipient ID (RID) of “0.”
5.	MID 5 is accepted, written to disk, and acknowledged.
6.	Receiving is successful and the receiving connection closes.
7.	Next the message delivery process starts. It is assigned a Delivery Connection ID (DCID) of “8” from 192.168.42.42 and to 10.5.3.25.
8.	The message delivery starts to RID “0.”
9.	Delivery is successful for MID 6 to RID “0.”
10.	The delivery connection closes.

Examples of Text Mail Log Entries

Following are some sample log entries based on various situations.

Message Injection and Delivery

A message is injected into the Email Security appliance for a single recipient. The message is successfully delivered.

```
Wed Jun 16 21:42:34 2004 Info: New SMTP ICID 282204970 interface mail.example.com
(1.2.3.4) address 2.3.4.5 reverse dns host unknown verified no
```

```
Wed Jun 16 21:42:34 2004 Info: ICID 282204970 SBRS None
```

```
Wed Jun 16 21:42:35 2004 Info: Start MID 200257070 ICID 282204970
```

```
Wed Jun 16 21:42:35 2004 Info: MID 200257070 ICID 282204970 From: <someone@foo.com>
```

```
Wed Jun 16 21:42:36 2004 Info: MID 200257070 ICID 282204970 RID 0 To: <user@example.com>
```

```
Wed Jun 16 21:42:38 2004 Info: MID 200257070 Message-ID
'<37gva9$5uvbhe@mail.example.com>'
```

```
Wed Jun 16 21:42:38 2004 Info: MID 200257070 Subject 'Hello'
```

```
Wed Jun 16 21:42:38 2004 Info: MID 200257070 ready 24663 bytes from <someone@foo.com>
```

```
Wed Jun 16 21:42:38 2004 Info: MID 200257070 antivirus negative
```

```
Wed Jun 16 21:42:38 2004 Info: MID 200257070 queued for delivery
```

```
Wed Jun 16 21:42:38 2004 Info: New SMTP DCID 2386069 interface 1.2.3.4 address 1.2.3.4

Wed Jun 16 21:42:38 2004 Info: Delivery start DCID 2386069 MID 200257070 to RID [0]

Wed Jun 16 21:42:38 2004 Info: ICID 282204970 close

Wed Jun 16 21:42:38 2004 Info: Message done DCID 2386069 MID 200257070 to RID [0]
[('X-SBRS', 'None')]

Wed Jun 16 21:42:38 2004 Info: MID 200257070 RID [0] Response 2.6.0
<37gva9$5uvbhe@mail.example.com> Queued mail for delivery

Wed Jun 16 21:42:43 2004 Info: DCID 2386069 close
```

Successful Message Delivery

```
Mon Mar 31 20:10:58 2003 Info: New SMTP DCID 5 interface 172.19.0.11 address
63.251.108.110

Mon Mar 31 20:10:58 2003 Info: Delivery start DCID 5 MID 4 to RID [0]

Mon Mar 31 20:10:58 2003 Info: Message done DCID 5 MID 4 to RID [0]

Mon Mar 31 20:11:03 2003 Info: DCID 5 close
```

Unsuccessful Message Delivery (Hard Bounce)

A message with two recipients is injected into the Email Security appliance. Upon delivery, the destination host returns a 5XX error, which indicates that the message cannot be delivered to either recipient. The appliance notifies the sender and removes the recipients from the queue.

```
Mon Mar 31 20:00:23 2003 Info: New SMTP DCID 3 interface 172.19.0.11 address
64.81.204.225

Mon Mar 31 20:00:23 2003 Info: Delivery start DCID 3 MID 4 to RID [0, 1]

Mon Mar 31 20:00:27 2003 Info: Bounced: DCID 3 MID 4 to RID 0 - 5.1.0 - Unknown address
error ('550', ['<george@yourdomain.com>... Relaying denied']) []

Mon Mar 31 20:00:27 2003 Info: Bounced: DCID 3 MID 4 to RID 1 - 5.1.0 - Unknown address
error ('550', ['<jane@yourdomain.com>... Relaying denied']) []

Mon Mar 31 20:00:32 2003 Info: DCID 3 close
```

Soft Bounce Followed by Successful Delivery

A message is injected into the Email Security appliance. On the first delivery attempt, the message soft bounces and is queued for future delivery. On the second attempt, the message is successfully delivered.

```
Mon Mar 31 20:10:58 2003 Info: New SMTP DCID 5 interface 172.19.0.11 address
63.251.108.110

Mon Mar 31 20:00:23 2003 Info: Delivery start DCID 3 MID 4 to RID [0, 1]
```



```
Mon Mar 31 20:00:23 2003 Info: Delayed: DCID 5 MID 4 to RID 0 - 4.1.0 - Unknown address
error ('466', ['Mailbox temporarily full.'])[]

Mon Mar 31 20:00:23 2003 Info: Message 4 to RID [0] pending till Mon Mar 31 20:01:23
2003

Mon Mar 31 20:01:28 2003 Info: DCID 5 close

Mon Mar 31 20:01:28 2003 Info: New SMTP DCID 16 interface PublicNet address 172.17.0.113

Mon Mar 31 20:01:28 2003 Info: Delivery start DCID 16 MID 4 to RID [0]

Mon Mar 31 20:01:28 2003 Info: Message done DCID 16 MID 4 to RID [0]

Mon Mar 31 20:01:33 2003 Info: DCID 16 close
```

Message Scanning Results for the scanconfig Command

You can use the `scanconfig` command to determine the system behavior when a message can not be deconstructed into its component parts (when removing attachments). The Options are `Deliver`, `Bounce`, or `Drop`.

The following example shows the Text Mail log with `scanconfig` set to `Deliver`.

```
Tue Aug 3 16:36:29 2004 Info: MID 256 ICID 44784 From: <test@virus.org>

Tue Aug 3 16:36:29 2004 Info: MID 256 ICID 44784 RID 0 To: <joe@example.com>

Tue Aug 3 16:36:29 2004 Info: MID 256 Message-ID '<137398.@virus.org>'

Tue Aug 3 16:36:29 2004 Info: MID 256 Subject 'Virus Scanner Test #22'

Tue Aug 3 16:36:29 2004 Info: MID 256 ready 1627 bytes from <test@virus.org>

Tue Aug 3 16:36:29 2004 Warning: MID 256, Message Scanning Problem: Continuation line
seen before first header

Tue Aug 3 16:36:29 2004 Info: ICID 44784 close

Tue Aug 3 16:36:29 2004 Info: MID 256 antivirus positive 'EICAR-AV-Test'

Tue Aug 3 16:36:29 2004 Info: Message aborted MID 256 Dropped by antivirus

Tue Aug 3 16:36:29 2004 Info: Message finished MID 256 done
```

The following example shows the Text Mail log with `scanconfig` set to `drop`.

```
Tue Aug 3 16:38:53 2004 Info: Start MID 257 ICID 44785

Tue Aug 3 16:38:53 2004 Info: MID 257 ICID 44785 From: test@virus.org

Tue Aug 3 16:38:53 2004 Info: MID 257 ICID 44785 RID 0 To: <joe@example.com>

Tue Aug 3 16:38:53 2004 Info: MID 257 Message-ID '<392912.@virus.org>'
```

```
Tue Aug 3 16:38:53 2004 Info: MID 25781 Subject 'Virus Scanner Test #22'

Tue Aug 3 16:38:53 2004 Info: MID 257 ready 1627 bytes from <test@virus.org>

Tue Aug 3 16:38:53 2004 Warning: MID 257, Message Scanning Problem: Continuation line
seen before first header

Tue Aug 3 16:38:53 2004 Info: Message aborted MID 25781 Dropped by filter 'drop_zip_c'

Tue Aug 3 16:38:53 2004 Info: Message finished MID 257 done

Tue Aug 3 16:38:53 2004 Info: ICID 44785 close
```

Message with Attachment

In this example, a content filter with condition “Message Body Contains” has been configured to enable identification of attachment names:

```
Sat Apr 23 05:05:42 2011 Info: New SMTP ICID 28 interface Management (192.0.2.10)
address 224.0.0.10 reverse dns host test.com verified yes

Sat Apr 23 05:05:42 2011 Info: ICID 28 ACCEPT SG UNKNOWNLIST match sbrs[-1.0:10.0]
SBRS 0.0

Sat Apr 23 05:05:42 2011 Info: Start MID 44 ICID 28

Sat Apr 23 05:05:42 2011 Info: MID 44 ICID 28 From: <sender1@example.com>

Sat Apr 23 05:05:42 2011 Info: MID 44 ICID 28 RID 0 To: <recipient1@example.org>

Sat Apr 23 05:05:42 2011 Info: MID 44 Message-ID '<000001cba32e$f24ff2e0$d6efd8a0$@com>'

Sat Apr 23 05:05:42 2011 Info: MID 44 Subject 'Message 001'

Sat Apr 23 05:05:42 2011 Info: MID 44 ready 240129 bytes from <sender1@example.com>

Sat Apr 23 05:05:42 2011 Info: MID 44 matched all recipients for per-recipient
policy DEFAULT in the inbound table

Sat Apr 23 05:05:42 2011 Info: ICID 28 close

Sat Apr 23 05:05:42 2011 Info: MID 44 interim verdict using engine: CASE
spam negative

Sat Apr 23 05:05:42 2011 Info: MID 44 using engine: CASE spam negative

Sat Apr 23 05:05:43 2011 Info: MID 44 attachment 'Banner.gif'

Sat Apr 23 05:05:43 2011 Info: MID 44 attachment '=D1=82=D0=B5=D1=81=D1=82.rst'

Sat Apr 23 05:05:43 2011 Info: MID 44 attachment 'Test=20Attachment.docx'

Sat Apr 23 05:05:43 2011 Info: MID 44 queued for delivery
```

Note that the second of the three attachments is Unicode. On terminals that cannot display Unicode, these attachments are represented in quoted-printable format.

Log Entries for Generated or Re-Written Messages

Some functions, such as rewrite/redirect actions (`alt-rcpt-to` filters, anti-spam rcpt rewrite, `bcc()` actions, anti-virus redirections, etc.), create new messages. When looking through the logs, you might need to check the results and add in further MIDs and possibly DCIDs. Entries such as these are possible:

```
Tue Jun  1 20:02:16 2004 Info: MID 14 generated based on MID 13 by bcc filter 'nonetest'
```

or:

```
Tue Jan  6 15:03:18 2004 Info: MID 2 rewritten to 3 by antispm
```

```
Fri May 14 20:44:43 2004 Info: MID 6 rewritten to 7 by alt-rcpt-to-filter filter  
'testfilt'
```

An interesting point to note about 'rewritten' entries is that they can appear after lines in the log indicating use of the new MID.

Messages Sent to the Spam Quarantine

When you send a message to the quarantine, the mail logs track the movement to and from the quarantine using the RCID (RPC connection ID) to identify the RPC connection. In the following mail log, a message is tagged as spam, and sent to the Spam Quarantine:

```
Wed Feb 14 12:11:40 2007 Info: Start MID 2317877 ICID 15726925
```

```
Wed Feb 14 12:11:40 2007 Info: MID 2317877 ICID 15726925 From: <HLD@chasehf.bfi0.com>
```

```
Wed Feb 14 12:11:40 2007 Info: MID 2317877 ICID 15726925 RID 0 To:  
<stevel@healthtrust.org>
```

```
Wed Feb 14 12:11:40 2007 Info: MID 2317877 Message-ID  
'<W1TH05606E5811BEA0734309D4BAF0.323.14460.pemailer44.DumpShot.2@email.chase.com>'
```

```
Wed Feb 14 12:11:40 2007 Info: MID 2317877 Subject 'Envision your dream home - Now make  
it a reality'
```

```
Wed Feb 14 12:11:40 2007 Info: MID 2317877 ready 15731 bytes from <HLD@chasehf.bfi0.com>
```

```
Wed Feb 14 12:11:40 2007 Info: MID 2317877 matched all recipients for per-recipient  
policy DEFAULT in the inbound table
```

```
Wed Feb 14 12:11:41 2007 Info: MID 2317877 using engine: CASE spam suspect
```

```
Wed Feb 14 12:11:41 2007 Info: EUQ: Tagging MID 2317877 for quarantine
```

```
Wed Feb 14 12:11:41 2007 Info: MID 2317877 antivirus negative
```

```
Wed Feb 14 12:11:41 2007 Info: MID 2317877 queued for delivery
```

Wed Feb 14 12:11:44 2007 Info: RPC Delivery start RCID 756814 MID 2317877 to local
IronPort Spam Quarantine

Wed Feb 14 12:11:45 2007 Info: EUQ: Quarantined MID 2317877

Wed Feb 14 12:11:45 2007 Info: RPC Message done RCID 756814 MID 2317877

Wed Feb 14 12:11:45 2007 Info: Message finished MID 2317877 done

Using Delivery Logs

Delivery logs record critical information about the email delivery operations of AsyncOS. The log messages are “stateless,” meaning that all associated information is recorded in each log message and users need not reference previous log messages for information about the current delivery attempt.

The delivery log records all information pertaining to email delivery operations for each recipient. All information is laid out in a logical manner and is human-readable after conversion using a utility provided by Cisco. The conversion tools are located at:

<http://support.ironport.com>

Delivery logs are recorded and transferred in a binary format for resource efficiency. Information recorded in delivery logs is shown in the following table:

Table 38-7 Delivery Log Statistics

Statistic	Description
Delivery status	Success (message was successfully delivered) or bounce (message was hard bounced)
Del_time	Delivery time
Inj_time	Injection time. $\text{del\_time} - \text{inj\_time}$ = time the recipient message stayed in the queue
Bytes	Message size
Mid	Message ID
Ip	Recipient host IP. The IP address of the host that received or bounced the recipient message
From	Envelope From, also known as Envelope Sender or MAIL FROM
Source_ip	Source host IP. The IP address of the host of the incoming message
Code	SMTP response code from recipient host
Reply	SMTP response message from recipient host
Rcpt Rid	Recipient ID. Recipient ID starts with <0>, messages with multiple recipients will have multiple recipient IDs
To	Envelope To
Attempts	Number of delivery attempts

If the delivery status was bounce, this additional information appears in the delivery log:

Table 38-8 **Delivery Log Bounce Information**

Statistic	Description
Reason	RFC 1893 Enhanced Mail Status Code interpretation of the SMTP response during the delivery
Code	SMTP response code from recipient host
Error	SMTP response message from recipient host

If you have set up logheaders (see [Logging Message Headers, page 38-42](#)), the header information appears after the delivery information:

Table 38-9 **Delivery Log Header Information**

Statistic	Description
Customer_data	XML tag marking the beginning of logged headers
Header Name	Name of the header
Value	Contents of the logged header

Examples of Delivery Log Entries

The examples in this section show a variety of Delivery Log entries.

Successful Message Delivery

```
<success del_time="Fri Jan 09 15:34:20.234 2004" inj_time="Fri Jan 09 15:33:38.623 2004"
bytes="202" mid="45949" ip="10.1.1.1" from="campaign1@yourdomain.com"
source_ip="192.168.102.1" code="250" reply="sent">

<rcpt rid="0" to="alsdfj.ajsdf1@alsdfj.d2.qa25.qa" attempts="1" />

</success>
```

Delivery Status Bounce

```
<bounce del_time="Sun Jan 05 08:28:33.073 2003" inj_time="Mon Jan 05 08:28:32.929 2003"
bytes="4074" mid="94157762" ip="0.0.0.0" from="campaign1@yourdomain.com"
source_ip="192.168.102.1" reason="5.1.0 - Unknown address error" code="550"
error=["Requested action not taken: mailbox unavailable"]">

<rcpt rid="0" to="user@sampldomain.com" attempts="1" />

</bounce>
```

Delivery Log Entry with Logheaders

```
<success del_time="Tue Jan 28 15:56:13.123 2003" inj_time="Tue Jan 28 15:55:17.696 2003"
bytes="139" mid="202" ip="10.1.1.13" from="campaign1@yourdomain.com"
source_ip="192.168.102.1" code="250" reply="sent">

<rcpt rid="0" to="user@sampldomain.com" attempts="1" />

<customer_data>

<header name="xname" value="sh"/>

</customer_data>

</success>
```

Using Bounce Logs

The bounce log records all information pertaining to each bounced recipient. Information recorded in bounce logs is shown in [Table 38-10](#).

Table 38-10 Bounce Log Statistics

Statistic	Description
Timestamp	The time of the bounce event
Log level	The level of detail in this bounce log
Bounce type	Bounced or delayed (for example, hard or soft-bounce)
MID/RID	Message ID and recipient ID
From	Envelope From

Table 38-10 Bounce Log Statistics (continued)

Statistic	Description
To	Envelope To
Reason	RFC 1893 Enhanced Mail Status Code interpretation of the SMTP response during the delivery
Response	SMTP response code and message from recipient host

In addition, if you have specified message size to log or setup logheaders (see [Logging Message Headers, page 38-42](#)), the message and header information will appear after the bounce information:

Table 38-11 Bounce Log Header Information

Header	The header name and content in the header
Message	Content of the message logged

Examples of Bounce Log Entries

Soft-Bounced Recipient (Bounce Type = Delayed)

```
Thu Dec 26 18:37:00 2003 Info: Delayed: 44451135:0
From:<campaign1@yourdomain.com> To:<user@sampledomain.com>
```

```
Reason: "4.1.0 - Unknown address error" Response: "('451',
['<user@sampledomain.com> Automated block triggered by suspicious
activity from your IP address (10.1.1.1). Have your system administrator
send e-mail to postmaster@sampledomain.com if you believe this block is
in error'])"
```

Hard-Bounced Recipient (Bounce Type = Bounced)

```
Thu Dec 26 18:36:59 2003 Info: Bounced: 45346670:0 From:<campaign1@yourdomain.com>
To:<user2@sampledomain.com>
```

```
Reason: "5.1.0 - Unknown address error" Response: "('550', ['There is no such active
account.'])"
```

Bounce Log with Message Body and Logheaders

```
Wed Jan 29 00:06:30 2003 Info: Bounced: 203:0 From:<campaign1@yourdomain.com>
To:<user@sampledomain.com>
```

```
Reason:"5.1.2 - Bad destination host" Response: "('000', [])" Headers: ['xname:
userID2333'] Message: Message-Id:
```

```
<lu5jak$6b@yourdomain.com>\015\012xname: userID2333\015\012subject:
Greetings.\015\012\015\012Hi Tom:'
```

**Note**

The text string \015\012 represents a line break (for example, CRLF).

Using Status Logs

Status logs record system statistics found in the CLI status commands, including `status`, `status detail`, and `dnsstatus`. The period of recording is set using the `setup` subcommand in `logconfig`. Each counter or rate reported in status logs is the value since the last time the counter was reset.

Reading Status Logs

Table 38-12 table shows the status log labels and the matching system statistics.

Table 38-12 Status Log Statistics

Statistic	Description
CPULd	CPU Utilization
DskIO	Disk I/O Utilization
RAMUtil	RAM Utilization
QKUsd	Queue Kilobytes Used
QKFre	Queue Kilobytes Free
CrtMID	Message ID (MID)
CrtICID	Injection Connection ID (ICID)
CRTDCID	Delivery Connection ID (DCID)
InjBytes	Total Injected Message Size in Bytes
InjMsg	Injected Messages
InjRcp	Injected Recipients
GenBncRcp	Generated Bounce Recipients
RejRcp	Rejected Recipients
DrpMsg	Dropped Messages
SftBncEvt	Soft Bounced Events
CmpRcp	Completed Recipients
HrdBncRcp	Hard Bounced Recipients
DnsHrdBnc	DNS Hard Bounces
5XXHrdBnc	5XX Hard Bounces
FltrHrdBnc	Filter Hard Bounces
ExpHrdBnc	Expired Hard Bounces
OtrHrdBnc	Other Hard Bounces
DlvRcp	Delivered Recipients
DelRcp	Deleted Recipients
GlbUnsbHt	Global Unsubscribe Hits
ActvRcp	Active Recipients
UnatmptRcp	Unattempted Recipients
AtmptRcp	Attempted Recipients
CrtCncln	Current Inbound Connections
CrtCncOut	Current Outbound Connections
DnsReq	DNS Requests
NetReq	Network Requests
CchHit	Cache Hits
CchMis	Cache Misses

Table 38-12 **Status Log Statistics (continued)**

Statistic	Description
CchEct	Cache Exceptions
CchExp	Cache Expired
CPUTm	Total CPU time used by the application
CPUETm	Elapsed time since the application started
MaxIO	Maximum disk I/O operations per second for the mail process
RamUsd	Allocated memory in bytes
SwIn	Memory swapped in.
SwOut	Memory swapped out.
SwPgIn	Memory paged in.
SwPgOut	Memory paged out.
MMLen	Total number of messages in the system
DstInMem	Number of destination objects in memory
ResCon	Resource conservation tarpit value. Acceptance of incoming mail is delayed by this number of seconds due to heavy system load
WorkQ	This is the number of messages currently in the work queue
QuarMsgs	Number of individual messages in policy, virus, or outbreak quarantine (messages present in multiple quarantines are counted only once)
QuarQKUsd	KBytes used by policy, virus, and outbreak quarantine messages
LogUsd	Percent of log partition used
AVLd	Percent CPU used by anti-virus scanning
CmrkLd	Percent CPU used by Cloudmark anti-spam scanning
SophLd	Percent CPU used by Sophos anti-spam scanning
McafLd	Percent CPU used by McAfee anti-virus scanning
CASELd	Percent CPU used by CASE scanning
TotalLd	Total CPU consumption
LogAvail	Amount of disk space available for log files
EuQ	Estimated number of messages in the Spam quarantine
EuQRls	Estimated number of messages in the Spam quarantine release queue

Status Log Example

```
Fri Feb 24 15:14:39 2006 Info: Status: CPULd 0 DskIO 0 RAMUtil 2 QKUsd 0 QKFre 8388608
CrtMID 19036 CrtICID 35284 CrtDCID 4861 InjMsg 13889 InjRcp 14230 GenBncRcp 12 RejRcp
6318 DrpMsg 7437 SftBncEvnt 1816 CmpRcp 6813 HrdBncRcp 18 DnsHrdBnc 2 5XXHrdBnc 15
FltrHrdBnc 0 ExpHrdBnc 1 OtrHrdBnc 0 DlvRcp 6793 DelRcp 2 GlbUnsbHt 0 ActvRcp 0
UnatmptRcp 0 AtmptRcp 0 CrtCncIn 0 CrtCncOut 0 DnsReq 143736 NetReq 224227 CchHit 469058
CchMis 504791 CchEct 15395 CchExp 55085 CPUTTm 228 CPUETm 181380 MaxIO 350 RAMUsd
21528056 MMLen 0 DstInMem 4 ResCon 0 WorkQ 0 QuarMsgs 0 QuarQKUsd 0 LogUsd 3 AVLd 0 BMLd
0 CASELd 3 TotalLd 3 LogAvail 17G EuQ 0 EuqRls 0
```

Using Domain Debug Logs

Domain debug logs record the client and server communication during an SMTP conversation between the Email Security appliance and a specified recipient host. This log type is primarily used to debug issues with specific recipient hosts.

Table 38-13 Domain Debug Log Statistics

Statistic	Description
Timestamp	The time of the bounce event
Log level	The level of detail in this bounce log
From	Envelope From
To	Envelope To
Reason	RFC 1893 Enhanced Mail Status Code interpretation of the SMTP response during the delivery
Response	SMTP response code and message from recipient host

Domain Debug Log Example

```
Sat Dec 21 02:37:22 2003 Info: 102503993 Sent: 'MAIL FROM:<daily@dailyf-y-i.net>'

Sat Dec 21 02:37:23 2003 Info: 102503993 Rcvd: '250 OK'

Sat Dec 21 02:37:23 2003 Info: 102503993 Sent: 'RCPT TO:<LLLSMILE@aol.com>'

Sat Dec 21 02:37:23 2003 Info: 102503993 Rcvd: '250 OK'

Sat Dec 21 02:37:23 2003 Info: 102503993 Sent: 'DATA'

Sat Dec 21 02:37:24 2003 Info: 102503993 Rcvd: '354 START MAIL INPUT, END WITH "." ON A
LINE BY ITSELF'

Sat Dec 21 02:37:24 2003 Info: 102503993 Rcvd: '250 OK'
```

Using Injection Debug Logs

Injection debug logs record the SMTP conversation between the Email Security appliance and a specified host connecting to the system. Injection debug logs are useful for troubleshooting communication problems between the Email Security appliance and a client initiating a connection from the Internet. The log records all bytes transmitted between the two systems and classifies them as “Sent to” the connecting host or “Received from” the connecting host.

You must designate the host conversations to record by specifying an IP address, an IP range, hostname, or partial hostname. Any connecting IP address within an IP range will be recorded. Any host within a partial domain will be recorded. The system performs reverse DNS lookups on connecting IP addresses to convert to hostnames. IP addresses without a corresponding PTR record in DNS will not match hostnames.

You must also specify the number of sessions to record.

Each line within an Injection Debug log contains the following information in [Table 38-14](#).

Table 38-14 Injection Debug Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted
ICID	The Injection Connection ID is a unique identifier that can be tied to the same connection in other log subscriptions
Sent/Received	Lines marked with “Sent to” are the actual bytes sent to the connecting host. Lines marked with “Received from” are the actual bytes received from the connecting host
IP Address	IP address of the connecting host

Injection Debug Log Example

```

Wed Apr 2 14:30:04 2003 Info: 6216 Sent to '172.16.0.22': '220 postman.example.com
ESMTP\015\012'

Wed Apr 2 14:30:04 2003 Info: 6216 Rcvd from '172.16.0.22': 'HELO
mail.remotehost.com\015\012'

Wed Apr 2 14:30:04 2003 Info: 6216 Sent to '172.16.0.22': '250
postman.example.com\015\012'

Wed Apr 2 14:30:04 2003 Info: 6216 Rcvd from '172.16.0.22': 'MAIL
FROM:<sender@remotehost.com>\015\012'

Wed Apr 2 14:30:04 2003 Info: 6216 Sent to '172.16.0.22': '250 sender
<sender@remotehost.com> ok\015\012'

Wed Apr 2 14:30:04 2003 Info: 6216 Rcvd from '172.16.0.22': 'RCPT
TO:<recipient@example.com>\015\012'

Wed Apr 2 14:30:04 2003 Info: 6216 Sent to '172.16.0.22': '250 recipient
<recipient@example.com> ok\015\012'

Wed Apr 2 14:30:04 Info: 6216 Rcvd from '172.16.0.22': 'DATA\015\012'

Wed Apr 2 14:30:04 2003 Info: 6216 Sent to '172.16.0.22': '354 go ahead\015\012'

Wed Apr 2 14:30:04 2003 Info: 6216 Rcvd from '172.16.0.22': 'To:
recipient@example.com\015\012Date: Apr 02 2003 10:09:44\015\012Subject: Test
Subject\015\012From: Sender <sender@remotehost.com>\015\012'

Wed Apr 2 14:30:04 2003 Info: 6216 Rcvd from '172.16.0.22': 'This is the content of the
message'

Wed Apr 2 14:30:04 Info: 6216 Sent to '172.16.0.22': '250 ok\015\012'

Wed Apr 2 14:30:04 Info: 6216 Rcvd from '172.16.0.22': 'QUIT\015\012'

Wed Apr 2 14:30:04 2003 Info: 6216 Sent to '172.16.0.22': '221
postman.example.com\015\012'

```

Using System Logs

Table 38-15 *System Log Statistics*

Statistic	Description
Timestamp	Time that the bytes were transmitted
Message	The logged event

System Log Example

In this example, the System log shows some commit entries, including the name of the user issuing the commit and the comment entered.

```

Wed Sep  8 18:02:45 2004 Info: Version: 4.0.0-206 SN: XXXXXXXXXXXX-XXX

Wed Sep  8 18:02:45 2004 Info: Time offset from UTC: 0 seconds

Wed Sep  8 18:02:45 2004 Info: System is coming up

Wed Sep  8 18:02:49 2004 Info: bootstrapping DNS cache

Wed Sep  8 18:02:49 2004 Info: DNS cache bootstrapped

Wed Sep  8 18:13:30 2004 Info: PID 608: User admin commit changes: SSW:Password

Wed Sep  8 18:17:23 2004 Info: PID 608: User admin commit changes: Completed Web::SSW

Thu Sep  9 08:49:27 2004 Info: Time offset from UTC: -25200 seconds

Thu Sep  9 08:49:27 2004 Info: PID 1237: User admin commit changes: Added a second CLI
log for examples

Thu Sep  9 08:51:53 2004 Info: PID 1237: User admin commit changes: Removed example CLI
log.

```

Using CLI Audit Logs

Table 38-16 CLI Audit Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted
PID	Process ID for the particular CLI session in which the command was entered
Message	The message consists of the CLI command entered, the CLI output (including menus, lists, etc.), and the prompt that is displayed

CLI Audit Log Example

In this example, the CLI Audit log shows that, for PID 16434, the following CLI commands were entered: `who`, `textconfig`.

```

Thu Sep  9 14:35:55 2004 Info: PID 16434: User admin entered 'who'; prompt was
'\nmail3.example.com> '

```

```

Thu Sep  9 14:37:12 2004 Info: PID 16434: User admin entered 'textconfig'; prompt was
'\nUsername Login Time Idle Time Remote Host What\n=====
===== \nadmin Wed 11AM 3m 45s 10.1.3.14 tail\nadmin 02:32PM
0s 10.1.3.14 cli\nmail3.example.com> '

```

```

Thu Sep  9 14:37:18 2004 Info: PID 16434: User admin entered ''; prompt was '\nThere are
no text resources currently defined.\n\nChoose the operation you want to perform:\n-
NEW - Create a new text resource.\n- IMPORT - Import a text resource from a file.\n[]> '

```

Using FTP Server Logs

Table 38-17 *FTP Server Log Statistics*

Statistic	Description
Timestamp	Time that the bytes were transmitted
ID	Connection ID. A separate ID for each FTP connection
Message	The message section of the log entry can be logfile status information, or FTP connection information (login, upload, download, logout, etc.)

FTP Server Log Example

In this example, the FTP Server log records a connection (ID:1). The IP address of the incoming connection is shown, as well as the activity (uploading and downloading files) and the logout.

```
Wed Sep  8 18:03:06 2004 Info: Begin Logfile

Wed Sep  8 18:03:06 2004 Info: Version: 4.0.0-206 SN: 00065BF3BA6D-9WFWC21

Wed Sep  8 18:03:06 2004 Info: Time offset from UTC: 0 seconds

Wed Sep  8 18:03:06 2004 Info: System is coming up

Fri Sep 10 08:07:32 2004 Info: Time offset from UTC: -25200 seconds

Fri Sep 10 08:07:32 2004 Info: ID:1 Connection from 10.1.3.14 on 172.19.0.86

Fri Sep 10 08:07:38 2004 Info: ID:1 User admin login SUCCESS

Fri Sep 10 08:08:46 2004 Info: ID:1 Upload wording.txt 20 bytes

Fri Sep 10 08:08:57 2004 Info: ID:1 Download words.txt 1191 bytes

Fri Sep 10 08:09:06 2004 Info: ID:1 User admin logout
```

Using HTTP Logs

Table 38-18 HTTP Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted
ID	Session ID
req	IP address of machine connecting
user	Username of user connecting
Message	Information regarding the actions performed. May include GET or POST commands or system status, etc.

HTTP Log Example

In this example, the HTTP log shows the admin user's interaction with the GUI (running the System Setup Wizard, etc.).

```

Wed Sep  8 18:17:23 2004 Info: http service on 192.168.0.1:80 redirecting to https port
443

Wed Sep  8 18:17:23 2004 Info: http service listening on 192.168.0.1:80

Wed Sep  8 18:17:23 2004 Info: https service listening on 192.168.0.1:443

Wed Sep  8 11:17:24 2004 Info: Time offset from UTC: -25200 seconds

Wed Sep  8 11:17:24 2004 Info: req:10.10.10.14 user:admin id:iaCkEh2h5rZknQarAecg POST
/system_administration/system_setup_wizard HTTP/1.1 303

Wed Sep  8 11:17:25 2004 Info: req:10.10.10.14 user:admin id:iaCkEh2h5rZknQarAecg GET
/system_administration/ssw_done HTTP/1.1 200

Wed Sep  8 11:18:45 2004 Info: req:10.10.10.14 user:admin id:iaCkEh2h5rZknQarAecg GET
/monitor/incoming_mail_overview HTTP/1.1 200

Wed Sep  8 11:18:45 2004 Info: req:10.10.10.14 user:admin id:iaCkEh2h5rZknQarAecg GET
/monitor/mail_flow_graph?injector=&width=365&interval=0&type=recipientsin&height=190
HTTP/1.1 200

Wed Sep  8 11:18:46 2004 Info: req:10.10.10.14 user:admin id:iaCkEh2h5rZknQarAecg GET
/monitor/classification_graph?injector=&width=325&interval=0&type=recipientsin&height=19
0 HTTP/1.1 200

Wed Sep  8 11:18:49 2004 Info: req:10.10.10.14 user:admin id:iaCkEh2h5rZknQarAecg GET
/monitor/quarantines HTTP/1.1 200

```


Using NTP Logs

Table 38-19 NTP Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted
Message	The message consists of either a Simple Network Time Protocol (SNTP) query to the server, or an adjust: message

NTP Log Example

In this example, the NTP log shows the appliance polling the NTP host twice.

```
Thu Sep 9 07:36:39 2004 Info: sntp query host 10.1.1.23 delay 653 offset -652
```

```
Thu Sep 9 07:36:39 2004 Info: adjust: time_const: 8 offset: -652us next_poll: 4096
```

```
Thu Sep 9 08:44:59 2004 Info: sntp query host 10.1.1.23 delay 642 offset -1152
```

```
Thu Sep 9 08:44:59 2004 Info: adjust: time_const: 8 offset: -1152us next_poll: 4096
```

Using Scanning Logs

The scanning log contains all LOG and COMMON messages for the appliance's scanning engines. See the Alerts section of the "System Administration" chapter for a list of available the COMMON and LOG alert messages.

Table 38-20 Scanning Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted
Message	The message consists of an application fault, sent alert, failed alert, or log error message for one of the scanning engines.

Scanning Log Example

In this example, the log shows the history of an appliance sending a warning alert concerning Sophos anti-virus.

```
Wed Feb 23 22:05:48 2011 Info: Internal SMTP system attempting to send a message to alerts@example.com with subject 'Warning <Anti-Virus> mail3.example.com: sophos antivirus - The Anti-Virus database on this system is...' (attempt #0).
```

```
Wed Feb 23 22:05:48 2011 Info: Internal SMTP system successfully sent a message to alerts@example.com with subject 'Warning <Anti-Virus> mail3.example.com: sophos antivirus - The Anti-Virus database on this system is...'.
```

```
Wed Feb 23 22:05:48 2011 Info: A Anti-Virus/Warning alert was sent to alerts@example.com with subject "Warning <Anti-Virus> mail3.example.com: sophos antivirus - The Anti-Virus database on this system is...".
```

Using Anti-Spam Logs

Table 38-21 Anti-Spam Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted
Message	The message consists of the check for the anti-spam updates, as well as the results (whether an update of the engine or the anti-spam rules was needed, etc.)

Anti-Spam Log Example

In this example, the anti-spam log shows the anti-spam engine checking for updates to spam definitions and CASE updates:

```
Fri Apr 13 18:59:47 2007 Info: case antispam - engine (19103) : case-daemon: server
successfully spawned child process, pid 19111
```

```
Fri Apr 13 18:59:47 2007 Info: case antispam - engine (19111) : startup: Region profile:
Using profile global
```

```
Fri Apr 13 18:59:59 2007 Info: case antispam - engine (19111) : fuzzy: Fuzzy plugin v7
successfully loaded, ready to roll
```

```
Fri Apr 13 19:00:01 2007 Info: case antispam - engine (19110) : uribllocal: running URI
blocklist local
```

```
Fri Apr 13 19:00:04 2007 Info: case antispam - engine (19111) : config: Finished loading
configuration
```

Using Anti-Virus Logs

Table 38-22 AntiVirus Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted
Message	The message consists of the check for the anti-virus update, as well as the results (whether an update of the engine or the virus definitions was needed, etc.)

Anti-Virus Log Example

In this example, the Anti-Virus log shows the Sophos anti-virus engine checking for updates to virus definitions (IDE) and the engine itself.

```
Thu Sep 9 14:18:04 2004 Info: Checking for Sophos Update
```

```
Thu Sep 9 14:18:04 2004 Info: Current SAV engine ver=3.84. No engine update needed
```

```
Thu Sep 9 14:18:04 2004 Info: Current IDE serial=2004090902. No update needed.
```

You can temporarily set this to DEBUG level to help diagnose why the anti-virus engine returns a particular verdict for a given message. The DEBUG logging information is verbose; use with caution.

Using Spam Quarantine Logs

Table 38-23 *Spam Log Statistics*

Statistic	Description
Timestamp	Time that the bytes were transmitted
Message	The message consists of actions taken (messages quarantined, released from quarantine, etc.).

Spam Quarantine Log Example

In this example, the log shows a message (MID 8298624) being released from the quarantine to admin@example.com.

```
Mon Aug 14 21:41:47 2006 Info: ISQ: Releasing MID [8298624, 8298625] for all

Mon Aug 14 21:41:47 2006 Info: ISQ: Delivering released MID 8298624 (skipping work
queue)

Mon Aug 14 21:41:47 2006 Info: ISQ: Released MID 8298624 to admin@example.com

Mon Aug 14 21:41:47 2006 Info: ISQ: Delivering released MID 8298625 (skipping work
queue)

Mon Aug 14 21:41:47 2006 Info: ISQ: Released MID8298625 to admin@example.com
```

Using Spam Quarantine GUI Logs

Table 38-24 *Spam GUI Log Statistics*

Statistic	Description
Timestamp	Time that the bytes were transmitted
Message	The message consists of actions taken, including user authentication, etc.

Spam Quarantine GUI Log Example

In this example, the log shows a successful authentication, login and logout:

```
Fri Aug 11 22:05:28 2006 Info: ISQ: Serving HTTP on 192.168.0.1, port 82

Fri Aug 11 22:05:29 2006 Info: ISQ: Serving HTTPS on 192.168.0.1, port 83

Fri Aug 11 22:08:35 2006 Info: Authentication OK, user admin

Fri Aug 11 22:08:35 2006 Info: logout:- user:pqufOtL6vyI5StCqhCfO session:10.251.23.228
```

```
Fri Aug 11 22:08:35 2006 Info: login:admin user:pquf0tL6vyI5StCqhCf0
session:10.251.23.228
```

```
Fri Aug 11 22:08:44 2006 Info: Authentication OK, user admin
```

Using LDAP Debug Logs

Table 38-25 LDAP Debug Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted
Message	LDAP Debug message

LDAP Debug Log Example



Note

Individual lines in log files are NOT numbered. They are numbered here only for sample purposes

```
1 Thu Sep 9 12:24:56 2004 Begin Logfile
2 Thu Sep 9 12:25:02 2004 LDAP: Masquerade query sun.masquerade address
employee@routing.qa to employee@mail.qa
3 Thu Sep 9 12:25:02 2004 LDAP: Masquerade query sun.masquerade address
employee@routing.qa to employee@mail.qa
4 Thu Sep 9 12:25:02 2004 LDAP: Masquerade query sun.masquerade address
employee@routing.qa to employee@mail.qa
5 Thu Sep 9 12:28:08 2004 LDAP: Clearing LDAP cache
6 Thu Sep 9 13:00:09 2004 LDAP: Query '(&(ObjectClass={g})(mailLocalAddress={a}))'
to server sun (sun.qa:389)
7 Thu Sep 9 13:00:09 2004 LDAP: After substitute, query is
'(&(ObjectClass=inetLocalMailRecipient)(mailLocalAddress=rroute.d00002b.loc@ldap.r
oute.local.add00002.qa))'
8 Thu Sep 9 13:00:09 2004 LDAP: connecting to server
9 Thu Sep 9 13:00:09 2004 LDAP: connected
10 Thu Sep 9 13:00:09 2004 LDAP: Query
(&(ObjectClass=inetLocalMailRecipient)(mailLocalAddress=rroute.d00002b.loc@ldap.ro
ute.local.add00002.qa)) returned 1 results
11 Thu Sep 9 13:00:09 2004 LDAP: returning: [<LDAP:>]
```

Use as a guide to reading the preceding log file.

Table 38-26 *Detail of LDAP Debug Log Example*

Line Number	Description
1.	The log file is initialized.
2.	The listener is configured to use LDAP for masquerading, specifically with the LDAP query named "sun.masquerade."
3.	
4.	
	The address employee@routing.qa is looked up in the LDAP server, a match is found, and the resulting masquerade address is employee@mail.qa, which will be written to the message headers and/or the envelope from, depending on the masquerade configuration.
5.	The user has manually run <code>ldapflush</code> .
6.	A query is about to be sent to sun.qa, port 389. The query template is: <code>(&(ObjectClass={g})(mailLocalAddress={a}))</code> .
	The {g} will be replaced by the groupname specified in the calling filter, either a rcpt-to-group or mail-from-group rule.
	The {a} will be replaced by the address in question.
7.	Now the substitution (described previously) takes place, and this is what the query looks like before it is sent to the LDAP server.
8.	
9.	The connection to the server is not yet established, so make a connection.
10.	The data that is sent to the server.
11.	The result is an empty positive, meaning one record was returned, but since the query didn't ask for any fields, there is no data to report. These are used for both group and accept queries when the query checks to see if there is a match in the database.

Using Safelist/Blocklist Logs

Table 38-27 shows the statistics recorded in safelist/blocklist logs.

Table 38-27 *Safelist/Blocklist Log Statistics*

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of actions taken, including user authentication, and so forth.

Safelist/Blocklist Log Example

In this example, the safelist/blocklist log shows the appliance creating database snapshots every two hours. It also shows when senders were added to the database.

```
Fri Sep 28 14:22:33 2007 Info: Begin Logfile Fri Sep 28 14:22:33 2007 Info: Version:
6.0.0-425 SN: XXXXXXXXXXXX-XXX Fri Sep 28 14:22:33 2007 Info: Time offset from UTC:
10800 seconds Fri Sep 28 14:22:33 2007 Info: System is coming up.
```

```
Fri Sep 28 14:22:33 2007 Info: SLBL: The database snapshot has been created.
```

```
Fri Sep 28 16:22:34 2007 Info: SLBL: The database snapshot has been created.
```

```
Fri Sep 28 18:22:34 2007 Info: SLBL: The database snapshot has been created.
```

```
Fri Sep 28 20:22:34 2007 Info: SLBL: The database snapshot has been created.
```

```
Fri Sep 28 22:22:35 2007 Info: SLBL: The database snapshot has been created.
```

```
.....
```

```
Mon Oct 1 14:16:09 2007 Info: SLBL: The database snapshot has been created.
```

```
Mon Oct 1 14:37:39 2007 Info: SLBL: The database snapshot has been created.
```

```
Mon Oct 1 15:31:37 2007 Warning: SLBL: Adding senders to the database failed.
```

```
Mon Oct 1 15:32:31 2007 Warning: SLBL: Adding senders to the database failed.
```

```
Mon Oct 1 16:37:40 2007 Info: SLBL: The database snapshot has been created.
```

Using Reporting Logs

Table 38-28 shows the statistics recorded in reporting logs.

Table 38-28 Reporting Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of actions taken, including user authentication, and so forth.

Reporting Log Example

In this example, the Reporting log shows the appliance set at the information log level.

```
Wed Oct 3 13:39:53 2007 Info: Period minute using 0 (KB)
```

```
Wed Oct 3 13:39:53 2007 Info: Period month using 1328 (KB)
```

```
Wed Oct 3 13:40:02 2007 Info: Update 2 registered appliance at 2007-10-03-13-40
```

```
Wed Oct 3 13:40:53 2007 Info: Pages found in cache: 1304596 (99%). Not found: 1692
```

```
Wed Oct 3 13:40:53 2007 Info: Period hour using 36800 (KB)
```

```

Wed Oct  3 13:40:53 2007 Info: Period day using 2768 (KB)

Wed Oct  3 13:40:53 2007 Info: Period minute using 0 (KB)

Wed Oct  3 13:40:53 2007 Info: Period month using 1328 (KB)

Wed Oct  3 13:40:53 2007 Info: HELPER checkpointed in 0.00580507753533 seconds

Wed Oct  3 13:41:02 2007 Info: Update 2 registered appliance at 2007-10-03-13-41

Wed Oct  3 13:41:53 2007 Info: Pages found in cache: 1304704 (99%). Not found: 1692

Wed Oct  3 13:41:53 2007 Info: Period hour using 36800 (KB)

Wed Oct  3 13:41:53 2007 Info: Period day using 2768 (KB)

Wed Oct  3 13:41:53 2007 Info: Period minute using 0 (KB)

Wed Oct  3 13:41:53 2007 Info: Period month using 1328 (KB)

Wed Oct  3 13:42:03 2007 Info: Update 2 registered appliance at 2007-10-03-13-42

```

Using Reporting Query Logs

Table 38-29 shows the statistics recorded in reporting query logs.

Table 38-29 **Reporting Query Log Statistics**

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of actions taken, including user authentication, and so forth.

Reporting Query Log Example

In this example, the reporting query log shows the appliance running a daily outgoing email traffic query for the period from August 29 to October 10, 2007.

```

Tue Oct  2 11:30:02 2007 Info: Query: Closing interval handle 811804479.

Tue Oct  2 11:30:02 2007 Info: Query: Closing interval handle 811804480.

Tue Oct  2 11:30:02 2007 Info: Query: Closing query handle 302610228.

Tue Oct  2 11:30:02 2007 Info: Query: Merge query with handle 302610229 for
['MAIL_OUTGOING_TRAFFIC_SUMMARY.

DETECTED_SPAM', 'MAIL_OUTGOING_TRAFFIC_SUMMARY.DETECTED_VIRUS',
'MAIL_OUTGOING_TRAFFIC_SUMMARY.THREAT_CONTEN

T_FILTER', 'MAIL_OUTGOING_TRAFFIC_SUMMARY.TOTAL_CLEAN_RECIPIENTS',
'MAIL_OUTGOING_TRAFFIC_SUMMARY.TOTAL_RECI

```

```

PIENTS_PROCESSED'] for rollup period "day" with interval range 2007-08-29 to 2007-10-01
with key constraints

None sorting on ['MAIL_OUTGOING_TRAFFIC_SUMMARY.DETECTED_SPAM'] returning results from
0 to 2 sort_ascendin

g=False.

Tue Oct  2 11:30:02 2007 Info: Query: Closing query handle 302610229.

Tue Oct  2 11:30:02 2007 Info: Query: Merge query with handle 302610230 for
['MAIL_OUTGOING_TRAFFIC_SUMMARY.

TOTAL_HARD_BOUNCES', 'MAIL_OUTGOING_TRAFFIC_SUMMARY.TOTAL_RECIPIENTS_DELIVERED',
'MAIL_OUTGOING_TRAFFIC_SUMM

ARY.TOTAL_RECIPIENTS'] for rollup period "day" with interval range 2007-08-29 to
2007-10-01 with key constra

ints None sorting on ['MAIL_OUTGOING_TRAFFIC_SUMMARY.TOTAL_HARD_BOUNCES'] returning
results from 0 to 2 sort

_ascending=False.

Tue Oct  2 11:30:02 2007 Info: Query: Closing query handle 302610230.

```

Using Updater Logs

Table 38-30 Updater Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of system service update information, as well as AsyncOS checking for updates and the scheduled date and time of the next update.

Updater Log Example

In this example, the logs show the appliance being updated with new McAfee Anti-Virus definitions.

```

Fri Sep 19 11:07:51 2008 Info: Starting scheduled update

Fri Sep 19 11:07:52 2008 Info: Acquired server manifest, starting update 11

Fri Sep 19 11:07:52 2008 Info: Server manifest specified an update for mcafee

Fri Sep 19 11:07:52 2008 Info: mcafee was signalled to start a new update

Fri Sep 19 11:07:52 2008 Info: mcafee processing files from the server manifest

Fri Sep 19 11:07:52 2008 Info: mcafee started downloading files

Fri Sep 19 11:07:52 2008 Info: mcafee downloading remote file
"http://stage-updates.ironport.com/mcafee/dat/5388"

```



```
Fri Sep 19 11:07:52 2008 Info: Scheduled next update to occur at Fri Sep 19 11:12:52
2008

Fri Sep 19 11:08:12 2008 Info: mcafee started decrypting files

Fri Sep 19 11:08:12 2008 Info: mcafee decrypting file
"mcafee/dat/5388" with method "des3_cbc"

Fri Sep 19 11:08:17 2008 Info: mcafee started decompressing files

Fri Sep 19 11:08:17 2008 Info: mcafee started applying files

Fri Sep 19 11:08:17 2008 Info: mcafee applying file "mcafee/dat/5388"

Fri Sep 19 11:08:18 2008 Info: mcafee verifying applied files

Fri Sep 19 11:08:18 2008 Info: mcafee updating the client manifest

Fri Sep 19 11:08:18 2008 Info: mcafee update completed

Fri Sep 19 11:08:18 2008 Info: mcafee waiting for new updates

Fri Sep 19 11:12:52 2008 Info: Starting scheduled update

Fri Sep 19 11:12:52 2008 Info: Scheduled next update to occur at Fri Sep 19 11:17:52
2008

Fri Sep 19 11:17:52 2008 Info: Starting scheduled update

Fri Sep 19 11:17:52 2008 Info: Scheduled next update to occur at Fri Sep 19 11:22:52
2008
```

Understanding Tracking Logs

Tracking logs record information about the email operations of AsyncOS. The log messages are a subset of the messages recorded in the mail logs.

The tracking logs are used by the appliance's message tracking component to build the message tracking database. Because the log files are consumed in the process of building the database, the tracking logs are transient. The information in tracking logs is not designed to be read or analyzed by humans.

You can also view tracking information from multiple Email Security appliances using the Cisco Security Management appliance.

Using Authentication Logs

The authentication log records successful user logins and unsuccessful login attempts.

Table 38-31 **Authentication Log Statistics**

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of the username of a user who attempted to log in to the appliance and whether the user was authenticated successfully.

Authentication Log Example

In this example, the log shows the log in attempts by users “admin,” “joe,” and “dan.”

```
Wed Sep 17 15:16:25 2008 Info: Begin Logfile

Wed Sep 17 15:16:25 2008 Info: Version: 6.5.0-262 SN: XXXXXXX-XXXXX

Wed Sep 17 15:16:25 2008 Info: Time offset from UTC: 0 seconds

Wed Sep 17 15:18:21 2008 Info: User admin was authenticated successfully.

Wed Sep 17 16:26:17 2008 Info: User joe failed authentication.

Wed Sep 17 16:28:28 2008 Info: User joe was authenticated successfully.

Wed Sep 17 20:59:30 2008 Info: User admin was authenticated successfully.

Wed Sep 17 21:37:09 2008 Info: User dan failed authentication.
```

Using Configuration History Logs

A configuration history log consists of a configuration file with an additional section listing the name of the user, a description of where in the configuration the user made changes, and the comment the user entered when committing the change. Each time a user commits a change, a new log is created containing the configuration file after the change.

Configuration History Log Example

In this example, the configuration history log shows that the user (admin) added a guest user to the table that defines which local users are allowed to log in to the system.

```
<?xml version="1.0" encoding="ISO-8859-1"?>

<!DOCTYPE config SYSTEM "config.dtd">

<!--

XML generated by configuration change.

Change comment: added guest user
```

User: admin

Configuration are described as:

This table defines which local users are allowed to log into the system.

Product: Cisco IronPort M160 Messaging Gateway(tm) Appliance

Model Number: M160

Version: 6.7.0-231

Serial Number: 000000000ABC-D0000000

Number of CPUs: 1

Memory (GB): 4

Current Time: Thu Mar 26 05:34:36 2009

Feature "Cisco IronPort Centralized Configuration Manager": Quantity = 10, Time Remaining = "25 days"

Feature "Centralized Reporting": Quantity = 10, Time Remaining = "9 days"

Feature "Centralized Tracking": Quantity = 10, Time Remaining = "30 days"

Feature "Centralized Spam Quarantine": Quantity = 10, Time Remaining = "30 days"

Feature "Receiving": Quantity = 1, Time Remaining = "Perpetual"

-->

<config>

Log Subscriptions

- [Configuring Log Subscriptions, page 38-39](#)
- [Creating a Log Subscription in the GUI, page 38-40](#)
- [Configuring Global Settings for Logging, page 38-40](#)
- [Rolling Over Log Subscriptions, page 38-43](#)
- [Configuring Host Keys, page 38-47](#)

Configuring Log Subscriptions

Use the Log Subscriptions page on the System Administration menu (or the `logconfig` command in the CLI) to configure a log subscription. Log subscriptions create log files that store information about AsyncOS activity, including errors. A log subscription is either retrieved or delivered (pushed) to another computer. Generally, log subscriptions have the following attributes:

Table 38-32 Log File Attributes

Attribute	Description
Log type	Defines the type of information recorded and the format of the log subscription. See Table 38-1, “Log Types,” on page 2 for more information.
Name	Nickname for the log subscription to be used for your future reference.
Rollover by File Size	The maximum size the file can reach before rolling over.
Rollover by Time	Sets the time interval for file rollovers.
Log level	Sets the level of detail for each log subscription.
Retrieval method	Defines how the log subscription will be obtained from the Email Security appliance.
Log filename	Used for the physical name of the file when written to disk. If multiple Email Security appliances are being used, the log filename should be unique to identify the system that generated the log file.

Log Levels

Log levels determine the amount of information delivered in a log. Logs can have one of five levels of detail. A more detailed setting creates larger log files and puts more drain on system performance. More detailed settings include all the messages contained in less detailed settings, plus additional messages. As the level of detail increases, system performance decreases.

**Note**

Log levels may be selected for all mail log types.

Table 38-33 Log Levels

Log Level	Description
Critical	The least detailed setting. Only errors are logged. Using this setting will not allow you to monitor performance and other important activities; however, the log files will not reach their maximum size as quickly. This log level is equivalent to the syslog level “Alert.”
Warning	All errors and warnings created by the system. Using this setting will not allow you to monitor performance and other important activities. This log level is equivalent to the syslog level “Warning.”
Information	The information setting captures the second-by-second operations of the system. For example, connections opened or delivery attempts. The Information level is the recommended setting for logs. This log level is equivalent to the syslog level “Info.”

Table 38-33 Log Levels (continued)

Log Level	Description
Debug	Use the Debug log level when you are trying to discover the cause of an error. Use this setting temporarily, and then return to the default level. This log level is equivalent to the syslog level “Debug.”
Trace	The Trace log level is recommended only for developers. Using this level causes a serious degradation of system performance and is not recommended. This log level is equivalent to the syslog level “Debug.”

Creating a Log Subscription in the GUI

Procedure

-
- Step 1** Choose **System Administration > Log Subscriptions**.
 - Step 2** Click **Add Log Subscription**.
 - Step 3** Select a log type and enter the log name (for the log directory) as well as the name for the log file itself.
 - Step 4** Specify the maximum file size before AsyncOS rolls over the log file as well as a time interval between rollovers. See [Rolling Over Log Subscriptions, page 38-43](#) for more information on rolling over log files.
 - Step 5** Select the log level. The available options are Critical, Warning, Information, Debug, or Trace.
 - Step 6** Configure the log retrieval method.
 - Step 7** Submit and commit your changes.
-

Editing Log Subscriptions

Procedure

-
- Step 1** Choose **System Administration > Log Subscriptions**.
 - Step 2** Click the name of the log in the Log Settings column.
 - Step 3** Make changes to the log subscription.
 - Step 4** Submit and commit your changes.
-

Configuring Global Settings for Logging

The system periodically records system measurements within the Text Mail Logs and the Status Logs. Use the **Edit Settings** button in the Global Settings section of the System Administration > Log Subscriptions page (or the `logconfig -> setup` command in the CLI) to configure:

- System metrics frequency. This is the amount of time, in seconds, that the system waits between recording measurements.
- Whether to record the Message-ID headers.

- Whether to record the remote response status code.
- Whether to record the subject header of the original message.
- A list of headers that should be logged for each message.

All logs optionally include the following three pieces of data:

1. Message-ID

When this option is configured, every message will have its Message ID header logged, if it is available. Note that this Message-ID may have come from the received message or may have been generated by AsyncOS itself. For example:

```
Tue Apr 6 14:38:34 2004 Info: MID 1 Message-ID Message-ID-Content
```

2. Remote Response

When this option is configured, every message will have its remote response status code logged, if it is available. For example:

```
Tue Apr 6 14:38:34 2004 Info: MID 1 RID [0] Response 'queued as 9C8B425DA7'
```

The remote response string is the human-readable text received after the response to the DATA command during the delivery SMTP conversation. In this example, the remote response after the connection host issued the data command is “queued as 9C8B425DA7.”

[...]

```
250 ok hostname
```

```
250 Ok: queued as 9C8B425DA7
```

Whitespace, punctuation, (and in the case of the 250 response, the OK characters) are stripped from the beginning of the string. Only whitespace is stripped from the end of the string. For example, Email Security appliances, by default, respond to the DATA command with this string: 250 Ok: Message MID accepted. So, the string “Message MID accepted” would be logged if the remote host were another Email Security appliance.

3. Original Subject Header

When this option is enabled, the original subject header of each message is included in the log.

```
Tue May 31 09:20:27 2005 Info: Start MID 2 ICID 2
```

```
Tue May 31 09:20:27 2005 Info: MID 2 ICID 2 From: <mary@example.com>
```

```
Tue May 31 09:20:27 2005 Info: MID 2 ICID 2 RID 0 To: <joe@example.com>
```

```
Tue May 31 09:20:27 2005 Info: MID 2 Message-ID '<44e4n$2@example.com>'
```

```
Tue May 31 09:20:27 2005 Info: MID 2 Subject 'Monthly Reports Due'
```

Logging Message Headers

In some cases, it is necessary to record the presence and contents of a message's headers as they pass through the system. You specify the headers to record in the Log Subscriptions Global Settings page (or via the `logconfig -> logheaders` subcommand in the CLI). The Email Security appliance records the specified message headers in the Text Mail Logs, the Delivery Logs, and the Bounce Logs. If the header is present, the system records the name of the header and the value. If a header is not present, nothing is recorded in the logs.


Note

The system evaluates all headers that are present on a message, at any time during the processing of the message for recording, regardless of the headers specified for logging.


Note

The RFC for the SMTP protocol is located at <http://www.faqs.org/rfcs/rfc2821.html> and defines user-defined headers.


Note

If you have configured headers to log via the `logheaders` command, the header information appears after the delivery information:

Table 38-34 Log Headers

Header name	Name of the header
Value	Contents of the logged header

For example, specifying “date, x-subject” as headers to be logged will cause the following line to appear in the mail log:

```
Tue May 31 10:14:12 2005 Info: Message done DCID 0 MID 3 to RID [0] [('date', 'Tue, 31
May 2005 10:13:18 -0700'), ('x-subject', 'Logging this header')]
```

Configuring Global Settings for Logging Using the GUI

Procedure

- Step 1** Choose **System Administration > Log Subscriptions**.
- Step 2** Scroll down to the **Global Settings** section.
- Step 3** Click **Edit Settings**.
- Step 4** Specify information including the system measurement frequency, whether to include Message-ID headers in mail logs, whether to include the remote response, and whether to include the original subject header of each message.
- Step 5** Enter any other headers you wish to include in the logs.
- Step 6** Submit and commit your changes.

Rolling Over Log Subscriptions

To prevent log files on the appliance from becoming too large, AsyncOS performs a “rollover” and archives a log file when it reaches a user-specified maximum file size or time interval and creates a new file for incoming log data. Based on the retrieval method defined for the log subscription, the older log file is stored on the appliance for retrieval or delivered to an external computer. See [Log Retrieval Methods, page 38-6](#) for more information on how to retrieve log files from the appliance.

When AsyncOS rolls over a log file, it performs the following actions:

- Renames the current log file with the timestamp of the rollover and a letter “s” extension signifying saved.
- Creates a new log file and designates the file as current with the “current” extension.
- Transfers the newly saved log file to a remote host (if using the push-based retrieval method).
- Transfers any previously unsuccessful log files from the same subscription (if using the push-based retrieval method).
- Deletes the oldest file in the log subscription if the total number of files to keep on hand has been exceeded (if using the poll-based retrieval method).

You define a log subscription’s rollover settings when creating or editing the subscription using the System Administration > Log Subscriptions page in the GUI or the `logconfig` command in the CLI. The two settings available for triggering a log file rollover are:

- A maximum file size.
- A time interval.

Rollover By File Size

AsyncOS rolls over log files when they reach a maximum file size to prevent them from using too much disk space. When defining a maximum file size for rollovers, use the suffix `m` for megabytes and `k` for kilobytes. For example, enter `10m` if you want AsyncOS to roll over the log file when it reaches 10 megabytes.

Rollover By Time

If you want to schedule rollovers to occur on a regular basis, you can select one of the following time intervals:

- **None.** AsyncOS only performs a rollover when the log file reaches the maximum file size.
- **Custom Time Interval.** AsyncOS performs a rollover after a specified amount of time has passed since the previous rollover. To create a custom time interval for scheduled rollovers, enter the number of days, hours, and minutes between rollovers using `d`, `h`, and `m` as suffixes.
- **Daily Rollover.** AsyncOS performs a rollover every day at a specified time. If you choose a daily rollover, enter the time of day you want AsyncOS to perform the rollover using the 24-hour format (HH:MM).

Only the GUI offers the Daily Rollover option. If you want to configure a daily rollover using the `logconfig` command in the CLI, choose the Weekly Rollover option and use an asterisk (*) to specify that AsyncOS should perform the rollover on every day of the week.

- **Weekly Rollover.** AsyncOS performs a rollover on one or more days of the week at a specified time. For example, you can set up AsyncOS to rollover the log file every Wednesday and Friday at midnight. To configure a weekly rollover, choose the days of the week to perform the rollover and the time of day in the 24-hour format (HH:MM).

If you are using the CLI, you can use a dash (-) to specify a range of days, an asterisk (*) to specify every day of the week, or a comma (,) to separate multiple days and times.

Table 38-35 shows how to use the CLI to roll over the files for a log subscription on Wednesday and Friday at midnight (00:00).

Table 38-35 **Weekly Log Rollover Settings in the CLI**

```
Do you want to configure time-based log files rollover? [N]> y
```

```
Configure log rollover settings:
```

```
1. Custom time interval.
```

```
2. Weekly rollover.
```

```
[1]> 2
```

```
1. Monday
```

```
2. Tuesday
```

```
3. Wednesday
```

```
4. Thursday
```

```
5. Friday
```

```
6. Saturday
```

```
7. Sunday
```

```
Choose the day of week to roll over the log files. Separate multiple days with comma,
or use "*" to specify every day of a week. Also you can use dash to specify a range
like "1-5":
```

```
[> 3, 5
```

```
Enter the time of day to rollover log files in 24-hour format (HH:MM). You can specify
hour as "*" to match every hour, the same for minutes. Separate multiple times of day
with comma:
```

```
[> 00:00
```

Rolling Over Log Subscriptions on Demand

To roll over log subscriptions immediately using the GUI:

Procedure

-
- Step 1** On the System Administration > Log Subscriptions page, mark the checkbox to the right of the logs you wish to roll over.
 - Step 2** Optionally, you can select all logs for rollover by marking the All checkbox.
 - Step 3** Once one or more logs have been selected for rollover, the **Rollover Now** button is enabled. Click the **Rollover Now** button to roll over the selected logs.
-

Viewing Recent Log Entries in the GUI

Before You Begin

You must have the HTTP or HTTPS service enabled on the Management interface in order to view logs via the GUI.

Procedure

-
- Step 1** Select **System Administration > Log Subscriptions**.
 - Step 2** Select the log subscription in the **Log Files** column of the table.
 - Step 3** Sign in.
 - Step 4** Select a log file to view it in your browser or to save it to disk.
-

Viewing Recent Log Entries in the CLI (tail Command)

AsyncOS supports a `tail` command, which shows the latest entries of configured logs on the appliance. Issue the `tail` command and select the number of a currently configured log to view it. Use Ctrl-C to exit from the `tail` command.

Example

In the following example, the `tail` command is used to view the system log. (This log tracks user comments from the `commit` command, among other things.) The `tail` command also accepts the name of a log to view as a parameter: `tail mail_logs`.

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

```
Currently configured logs:
```

1. "antispam" Type: "Anti-Spam Logs" Retrieval: Manual Download
2. "antivirus" Type: "Anti-Virus Logs" Retrieval: Manual Download
3. "asarchive" Type: "Anti-Spam Archive" Retrieval: Manual Download
4. "authentication" Type: "Authentication Logs" Retrieval: Manual Download
5. "avarchive" Type: "Anti-Virus Archive" Retrieval: Manual Download
6. "bounces" Type: "Bounce Logs" Retrieval: Manual Download
7. "cli_logs" Type: "CLI Audit Logs" Retrieval: Manual Download
8. "encryption" Type: "Encryption Logs" Retrieval: Manual Download
9. "error_logs" Type: "IronPort Text Mail Logs" Retrieval: Manual Download
10. "euq_logs" Type: "IronPort Spam Quarantine Logs" Retrieval: Manual Download
11. "euggui_logs" Type: "IronPort Spam Quarantine GUI Logs" Retrieval: Manual Download
12. "ftpd_logs" Type: "FTP Server Logs" Retrieval: Manual Download
13. "gui_logs" Type: "HTTP Logs" Retrieval: Manual Download
14. "mail_logs" Type: "IronPort Text Mail Logs" Retrieval: Manual Download
15. "reportd_logs" Type: "Reporting Logs" Retrieval: Manual Download
16. "reportqueryd_logs" Type: "Reporting Query Logs" Retrieval: Manual Download
17. "scanning" Type: "Scanning Logs" Retrieval: Manual Download
18. "slbld_logs" Type: "Safe/Block Lists Logs" Retrieval: Manual Download
19. "sntpd_logs" Type: "NTP logs" Retrieval: Manual Download
20. "status" Type: "Status Logs" Retrieval: Manual Download
21. "system_logs" Type: "System Logs" Retrieval: Manual Download
22. "trackerd_logs" Type: "Tracking Logs" Retrieval: Manual Download
23. "updater_logs" Type: "Updater Logs" Retrieval: Manual Download

Enter the number of the log you wish to tail.

[> 19

Press Ctrl-C to stop.

```
Mon Feb 21 12:25:10 2011 Info: PID 274: User system commit changes: Automated Update for
Quarantine Delivery Host
```

```
Mon Feb 21 23:18:10 2011 Info: PID 19626: User admin commit changes:
```

```
Mon Feb 21 23:18:10 2011 Info: PID 274: User system commit changes: Updated filter logs
config
```

```
Mon Feb 21 23:46:06 2011 Info: PID 25696: User admin commit changes: Receiving
suspended.
```

```
^Cmail3.example.com>
```

Configuring Host Keys

Use the `logconfig -> hostkeyconfig` subcommand to manage host keys for use with SSH when pushing logs to other servers from the Email Security appliance. SSH servers must have a pair of host keys, one private and one public. The private host key resides on the SSH server and cannot be read by remote machines. The public host key is distributed to any client machine that needs to interact with the SSH server.



Note

To manage user keys, see [Managing Secure Shell \(SSH\) Keys, page 32-28](#).

The `hostkeyconfig` subcommand performs the following functions:

Table 38-36 Managing Host Keys - List of Subcommands

Command	Description
New	Add a new key.
Edit	Modify an existing key.
Delete	Delete an existing key.
Scan	Automatically download a host key.
Print	Display a key.
Host	Display system host keys. This is the value to place in the remote system's 'known_hosts' file.
Fingerprint	Display system host key fingerprints.
User	Display the public key of the system account that pushes the logs to the remote machine. This is the same key that is displayed when setting up an SCP push subscription. This is the value to place in the remote system's 'authorized_keys' file.

In the following example, AsyncOS scans for host keys and add them for the host:

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

```
Currently configured logs:
```

```
[ list of logs ]
```

Choose the operation you want to perform:

- NEW - Create a new log.
- EDIT - Modify a log subscription.
- DELETE - Remove a log subscription.
- SETUP - General settings.
- LOGHEADERS - Configure headers to log.
- HOSTKEYCONFIG - Configure SSH host keys.

```
[> hostkeyconfig
```

Currently installed host keys:

```
1. mail3.example.com ssh-dss [ key displayed ]
```

Choose the operation you want to perform:

- NEW - Add a new key.
- EDIT - Modify a key.
- DELETE - Remove a key.
- SCAN - Automatically download a host key.
- PRINT - Display a key.
- HOST - Display system host keys.
- FINGERPRINT - Display system host key fingerprints.
- USER - Display system user keys.

```
[> scan
```

Please enter the host or IP address to lookup.

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

Choose the ssh protocol type:

```
1. SSH1:rsa
```

```
2. SSH2:rsa
```

```
3. SSH2:dsa
```

```
4. All
```

```
[4]>
```

```
SSH2:dsa
```

```
mail3.example.com ssh-dss
```

```
[ key displayed ]
```

```
SSH2:rsa
```

```
mail3.example.com ssh-rsa
```

```
[ key displayed ]
```

```
SSH1:rsa
```

```
mail3.example.com 1024 35
```

```
[ key displayed ]
```

```
Add the preceding host key(s) for mail3.example.com? [Y]>
```

```
Currently installed host keys:
```

```
1. mail3.example.com ssh-dss [ key displayed ]
```

```
2. mail3.example.com ssh-rsa [ key displayed ]
```

```
3. mail3.example.com 1024 35 [ key displayed ]
```

```
Choose the operation you want to perform:
```

```
- NEW - Add a new key.
```

```
- EDIT - Modify a key.
```

```
- DELETE - Remove a key.
```

- SCAN - Automatically download a host key.
- PRINT - Display a key.
- HOST - Display system host keys.
- FINGERPRINT - Display system host key fingerprints.
- USER - Display system user keys.

[]>

Currently configured logs:

[*list of configured logs*]

Choose the operation you want to perform:

- NEW - Create a new log.
- EDIT - Modify a log subscription.
- DELETE - Remove a log subscription.
- SETUP - General settings.
- LOGHEADERS - Configure headers to log.
- HOSTKEYCONFIG - Configure SSH host keys.

[]>



Centralized Management Using Clusters

- [Overview of Centralized Management Using Clusters, page 39-1](#)
- [Cluster Requirements, page 39-2](#)
- [Cluster Organization, page 39-2](#)
- [Creating and Joining a Cluster, page 39-4](#)
- [Managing Clusters, page 39-10](#)
- [Administering a Cluster from the GUI, page 39-15](#)
- [Cluster Communication, page 39-18](#)
- [Loading a Configuration in Clustered Appliances, page 39-22](#)
- [Best Practices and Frequently Asked Questions, page 39-24](#)

Overview of Centralized Management Using Clusters

The Cisco centralized management feature allows you to manage and configure multiple appliances at the same time, reducing administration time and ensuring a consistent configuration across your network. You do not need to purchase additional hardware for managing multiple appliances. The centralized management feature provides increased reliability, flexibility, and scalability within your network, allowing you to manage globally while complying with local policies.

A *cluster* is defined as a set of machines that share configuration information. Within the cluster, machines (Cisco appliances) are divided into *groups*; every cluster will contain at least one group. A given machine is a member of one and only one group. An administrator user can configure different elements of the system on a cluster-wide, group-wide, or per-machine basis, enabling the segmentation of Cisco appliances based on network, geography, business unit, or other logical relationships.

Clusters are implemented as a *peer-to-peer* architecture; there is no master/slave relationship within a cluster. You may log into any machine to control and administer the cluster. (Some configuration commands, however, are limited. See [Restricted Commands, page 39-14](#).)


The user database is shared across all machines in the cluster. That is, there will be only one set of users and one administrator user (with the associated passwords) for an entire cluster. All machines that join a cluster will share a single administrator password which is referred to as the *admin password* of the cluster.

Cluster Requirements

- Machines in a cluster must have resolvable hostnames in DNS. Alternatively, you can use IP addresses instead, but you may not mix the two.
See [DNS and Hostname Resolution, page 39-18](#). Cluster communication is normally initiated using the DNS hostnames of the machines.
- A cluster must consist entirely of machines running the same version of AsyncOS.
See [Upgrading Machines in a Cluster, page 39-12](#) for how to upgrade members of a cluster.
- Machines can either join the cluster via SSH (typically on port 22) *or* via the Cluster Communication Service (CCS).
See [Cluster Communication, page 39-18](#).
- Once machines have joined the cluster, they can communicate via SSH or via Cluster Communication Service. The port used is configurable. SSH is typically enabled on port 22, and by default CCS is on port 2222, but you can configure either of these services on a different port.
In addition to the normal firewall ports that must be opened for the appliance, clustered machines communicating via CCS must be able to connect with each other via the CCS port. See [Cluster Communication, page 39-18](#).
- You must use the Command Line Interface (CLI) command `clusterconfig` to create, join, or configure clusters of machines.
Once you have created a cluster, you can manage non-cluster configuration settings from either the GUI or the CLI.
See [Creating and Joining a Cluster, page 39-4](#) and [Administering a Cluster from the GUI, page 39-15](#).

Cluster Organization

Within a cluster, configuration information is divided into 3 groupings or *levels*. The top level describes cluster settings; the middle level describes group settings; and the lowest level describes machine-specific settings.

Figure 39-1 Cluster Level Hierarchy

Within each level there will be one or more specific members for which settings may be configured; these are referred to as *modes*. A mode refers to a named member at a specified level. For example, the group “usa” represents one of two group modes in the diagram. While levels are a general term, modes are specific; modes are always referred to by name. The cluster depicted in Figure 39-1 has six modes.

Although settings are configured at a given level, they are always configured **for** a specific mode. It is not necessary to configure settings for all modes within a level. The cluster mode is a special case. Because there can only be one cluster, all settings configured for the cluster mode can be said to be configured at the cluster level.

You should normally configure most settings at the cluster level. However, settings that have been specifically configured at lower levels will *override* settings configured at higher levels. Thus, you can override cluster-mode settings with group-mode or machine-mode settings.

For example, you might start by configuring the Good Neighbor Table in cluster mode; all machines in the cluster would use that configuration. Then, you might also configure this table in machine mode for machine *newyork*. In this case, all other machines in the cluster will still use the good neighbor table defined at the cluster level, but the machine *newyork* will override the cluster settings with its individual machine mode settings.

The ability to override cluster settings for specific groups or machines gives you a lot of flexibility. However, if you find yourself configuring many settings individually in machine mode, you will lose much of the ease of administration that clusters were intended to provide.

Initial Configuration Settings

For most features, when you begin to configure settings for a new mode, those settings will initially be empty by default. There is a distinction between empty settings and having no settings in a mode. As an example, consider a very simple cluster composed of one group and one machine. Imagine that you have an LDAP query configured at the cluster level. There are no settings configured at the group or machine levels:

Cluster	(ldap queries: a, b, c)
Group	
Machine	

Now, imagine that you create new LDAP query settings for the group. The result will be something like this:

Cluster	(ldap queries: a, b, c)
Group	(ldap queries: None)
Machine	

The group-level settings now override the cluster-level setting; however, the new group settings are initially empty. The group mode does not actually have any LDAP queries of its own configured. Note that a machine within this group will inherit this “empty” set of LDAP queries from the group.

Next, you can add an LDAP query to the group, for example:

Cluster	(ldap queries: a, b, c)
Group	(ldap queries: d)
Machine	

Now the cluster level has one set of queries configured while the group has another set of queries. The machine will inherit its queries from the group.

Creating and Joining a Cluster

You cannot create or join a cluster from the Graphical User Interface (GUI). You must use the Command Line Interface (CLI) to create, join, or configure clusters of machines. Once you have created a cluster, you can change configuration settings from either the GUI or the CLI.

The clusterconfig Command

A machine can create or join a cluster only via the `clusterconfig` command.

- When a new cluster is *created*, all of that cluster’s initial settings will be inherited from the machine that creates the cluster. If a machine was previously configured in “standalone” mode, its standalone settings are used when creating the cluster.
- When a machine *joins* a cluster, all of that machine’s clusterable settings will be inherited from the cluster level. In other words, everything except certain machine-specific settings (IP addresses, etc) will be lost and will be replaced with the settings from the cluster and/or the group selected for that machine to join. If a machine was previously configured in “standalone” mode, its standalone settings are used when creating the cluster, and no settings at the machine level are maintained.

If the current machine is not already part of a cluster, issuing the `clusterconfig` command presents the option to join an existing cluster or create a new one.

```
newyork.example.com> clusterconfig
```

```
Do you want to join or create a cluster?
```

```
1. No, configure as standalone.
```

2. Create a new cluster.
3. Join an existing cluster over SSH.
4. Join an existing cluster over CCS.

[1]> 2

Enter the name of the new cluster.

[> **americas**

New cluster committed: Wed Jun 22 10:02:04 2005 PDT

Creating a cluster takes effect immediately, there is no need to commit.

Cluster americas

Choose the operation you want to perform:

- ADDGROUP - Add a cluster group.
- SETGROUP - Set the group that machines are a member of.
- RENAMEGROUP - Rename a cluster group.
- DELETEGROUP - Remove a cluster group.
- REMOVEMACHINE - Remove a machine from the cluster.
- SETNAME - Set the cluster name.
- LIST - List the machines in the cluster.
- LISTDETAIL - List the machines in the cluster with detail.
- DISCONNECT - Temporarily detach machines from the cluster.
- RECONNECT - Restore connections with machines that were previously detached.
- PREPJOIN - Prepare the addition of a new machine over CCS.

[>

At this point you can add machines to the new cluster. Those machines can communicate via SSH or CCS.

Joining an Existing Cluster

From the host you want to add to the cluster, issue the `clusterconfig` command to join the existing cluster. You can choose to join the cluster over SSH or over CCS (cluster communication service).

In order to join a host to an existing cluster, you must:

- be able to validate the SSH host key of a machine in the cluster
- know the IP address of a machine in the cluster and be able to connect to this machine in the cluster (for example, via SSH or CCS)
- know the administrator password for the admin user on a machine belonging to the cluster

Joining an Existing Cluster over SSH

The following table demonstrates adding the machine `losangeles.example.com` to the cluster using the SSH option.

```
losangeles.example.com> clusterconfig
```

```
Do you want to join or create a cluster?
```

1. No, configure as standalone.
2. Create a new cluster.
3. Join an existing cluster over SSH.
4. Join an existing cluster over CCS.

```
[1]> 3
```

While joining a cluster, you will need to validate the SSH host key of the remote machine to which you are joining. To get the public host key

fingerprint of the remote host, connect to the cluster and run: `logconfig -> hostkeyconfig -> fingerprint`.

WARNING: All non-network settings will be lost. System will inherit the values set at the group or cluster mode for the non-network settings. Ensure that the cluster settings are compatible with your network settings (e.g. `dnsconfig` settings)

```
Do you want to enable the Cluster Communication Service on
```

```
losangeles.example.com? [N]> n
```

```
Enter the IP address of a machine in the cluster.
```

[1]> *IP address is entered*

Enter the remote port to connect to. The must be the normal admin ssh port, not the CCS port.

[22]> 22

Enter the admin password for the cluster.

The administrator password for the clustered machine is entered

Please verify the SSH host key for IP address:

Public host key fingerprint: xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx

Is this a valid key for this host? [Y]> **y**

Joining cluster group Main_Group.

Joining a cluster takes effect immediately, there is no need to commit.

Cluster americas

Choose the operation you want to perform:

- ADDGROUP - Add a cluster group.
- SETGROUP - Set the group that machines are a member of.
- RENAMEGROUP - Rename a cluster group.
- DELETEGROUP - Remove a cluster group.
- REMOVEMACHINE - Remove a machine from the cluster.
- SETNAME - Set the cluster name.
- LIST - List the machines in the cluster.
- LISTDETAIL - List the machines in the cluster with detail.
- DISCONNECT - Temporarily detach machines from the cluster.
- RECONNECT - Restore connections with machines that were previously detached.
- PREPJOIN - Prepare the addition of a new machine over CCS.

```
[ ]>  
  
(Cluster americas)>
```

Joining an Existing Cluster over CCS

Use CCS instead of SSH if you cannot use SSH. The only advantage of CCS is that only cluster communication happens over that port (no user logins, SCP, etc). To add another machine to an existing cluster via CCS, use the `prepjoin` subcommand of `clusterconfig` to prepare the machine to be added to the cluster. In this example, the `prepjoin` command is issued on the machine `newyork` to prepare the machine `losangeles` to be added to the cluster.

The `prepjoin` command involves obtaining the user key of the host you want to add to the cluster by typing `clusterconfig prepjoin print` in the CLI of that host, and then copying the key into the command line of the host that is currently in the cluster.

Choose the operation you want to perform:

- ADDGROUP - Add a cluster group.
- SETGROUP - Set the group that machines are a member of.
- RENAMEGROUP - Rename a cluster group.
- DELETEGROUP - Remove a cluster group.
- REMOVEMACHINE - Remove a machine from the cluster.
- SETNAME - Set the cluster name.
- LIST - List the machines in the cluster.
- LISTDETAIL - List the machines in the cluster with detail.
- DISCONNECT - Temporarily detach machines from the cluster.
- RECONNECT - Restore connections with machines that were previously detached.
- PREPJOIN - Prepare the addition of a new machine over CCS.

```
[ ]> prepjoin
```

Prepare Cluster Join Over CCS

No host entries waiting to be added to the cluster.

Choose the operation you want to perform:

- NEW - Add a new host that will join the cluster.

```
[> new
```

Enter the hostname of the system you want to add.

```
[> losangeles.example.com
```

Enter the serial number of the host mail3.example.com.

```
[> unique serial number is added
```

Enter the user key of the host losangeles.example.com. This can be obtained by typing "clusterconfig prepjoin print" in the CLI on mail3.example.com. Press enter on a blank line to finish.

unique user key from output of prepjoin print is pasted

Host losangeles.example.com added.

Prepare Cluster Join Over CCS

1. losangeles.example.com (serial-number)

Choose the operation you want to perform:

- NEW - Add a new host that will join the cluster.
- DELETE - Remove a host from the pending join list.

```
[>
```

Once a machine is already part of a cluster, the `clusterconfig` command allows you to configure various settings for the cluster.

```
(Cluster Americas)> clusterconfig
```

Cluster americas

Choose the operation you want to perform:

- ADDGROUP - Add a cluster group.

- SETGROUP - Set the group that machines are a member of.
- RENAMEGROUP - Rename a cluster group.
- DELETEGROUP - Remove a cluster group.
- REMOVEMACHINE - Remove a machine from the cluster.
- SETNAME - Set the cluster name.
- LIST - List the machines in the cluster.
- LISTDETAIL - List the machines in the cluster with detail.
- DISCONNECT - Temporarily detach machines from the cluster.
- RECONNECT - Restore connections with machines that were previously detached.
- PREPJOIN - Prepare the addition of a new machine over CCS.

[]>

Adding Groups

All clusters must contain at least one group. When you create a new cluster, a default group called `Main_Group` is created automatically. However, you may decide to create additional groups within your cluster. This example shows how to create additional groups within an existing cluster and assign machines to the new group(s).

Procedure

-
- Step 1** Issue the `clusterconfig` command.
 - Step 2** Choose the `addgroup` subcommand and enter the name of the new group.
 - Step 3** Use the `setgroup` subcommand to choose machines for the new group.

Managing Clusters

Administering a Cluster from the CLI

For machines that are part of a cluster, the CLI can be switched into different *modes*. Recall that a mode refers to a specific, named, member of a level.

The CLI mode determines precisely where a configuration setting will be modified. The default is “machine” mode for the machine the user logged into, the “login host.”

Use the `clustermode` command to switch between different modes.

Table 39-1 **Administering Clusters**

Command Example	Description
<code>clustermode</code>	Prompt to switch cluster mode
<code>clustermode group northamerica</code>	Switch to group mode for the group “northamerica”
<code>clustermode machine losangeles.example.com</code>	Switch to machine mode for the machine “losangeles”

The prompt in the CLI changes to indicate your current mode.

```
(Cluster Americas)>
```

or

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

In machine mode, the prompt will include the fully qualified domain name of the machine.

Copying and Moving Settings

All non-restricted (see [Restricted Commands, page 39-14](#)) commands have new operations: `CLUSTERSHOW` and `CLUSTERSET`. `CLUSTERSHOW` is used to show in which modes a command is configured (see [New Operation Added, page 39-14](#)). The `CLUSTERSET` operation allows you to move or copy the current settings (configurable with the current command) from one mode to another or between levels (e.g. from a machine to a group).

A *copy* retains the settings for the current mode. A *move* resets (clears) the configuration of the current mode; i.e., following a move, no settings will be configured for the current mode.

For example, if you have configured Good Neighbor Table settings (the `destconfig` command) for group `northamerica`, and you decide that you want the entire cluster to have these settings, you can use the `clusterset` operation from within the `destconfig` command to copy (or move) the current settings to the cluster mode. (See [Experimenting with New Configurations, page 39-11](#).)



Caution

Exercise caution when moving or copying configuration settings to avoid inconsistent dependencies. For example, if you move or copy listeners with disclaimer stamping configured to another machine, and that new machine does not have the same disclaimers configured, disclaimer stamping will not be enabled on the new machine.

Experimenting with New Configurations

One of the most advantageous ways to use clusters is to experiment with new configuration settings. First you make changes at the machine mode, in an isolated environment. Then, when you are satisfied with your configuration, you move those configuration changes up to the cluster mode to make them available on all machines.

The following example shows the steps to change a listener setting on one machine and then publish the setting to the rest of the cluster when ready. Because listeners are normally configured at the cluster level, the example starts by pulling the configuration down to machine mode on one machine before making and testing the changes. You should test experimental changes of this type on one machine before making the change to the other machines in the cluster.

Procedure

-
- Step 1** Use the `clustermode cluster` command to change to the cluster mode.
- Remember: the `clustermode` command is the CLI command you use to change modes to the cluster, group, and machine levels.
- Step 2** Type `listenerconfig` to see the listener settings configured for the cluster.
- Step 3** Choose the machine you want to experiment with, then use the `clusterset` command to copy settings from the cluster “down” to machine mode.
- Step 4** Use the `clustermode` command to navigate to machine mode for the experimental machine, e.g.:
- ```
clustermode machine newyork.example.com
```
- Step 5** In machine mode, on the experimental machine, issue the `listenerconfig` command to make changes specifically for the experimental machine.
- Step 6** Commit the changes.
- Step 7** Continue to experiment with the configuration changes on the experimental machine, remembering to commit the changes.
- Step 8** When you are ready to apply your new settings to all the other machines, use the `clusterset` command to move the settings up to the cluster mode.
- Step 9** Commit the changes.

## Leaving a Cluster Permanently (Removal)

You use the `REMOVEMACHINE` operation of `clusterconfig` to remove a machine permanently from a cluster. When a machine is permanently removed from a cluster, its configuration is “flattened” such that it will work the same as it did when it was part of the cluster. For example, if there is only a cluster-mode Global Unsubscribe table, the Global Unsubscribe table data will be copied to the machine’s local configuration when the machine is removed from the cluster.

## Upgrading Machines in a Cluster

A cluster does not allow the connected machines to have different versions of AsyncOS.

Before you install an AsyncOS upgrade, you need to disconnect each machine in the cluster via the `clusterconfig` command. After you upgrade all the machines, the cluster can be reconnected via the `clusterconfig` command. You can have two separate clusters running while you upgrade machines to the same version. You can also upgrade clustered machines on the GUI Upgrades page.

You can download the upgrade in the background so that you do not need to disconnect the cluster machines until you are ready to install the upgrade.

**Note**

If you use the `upgrade` command before disconnecting the individual machine from the cluster, AsyncOS disconnects all the machines in the cluster. Cisco Systems recommends that you disconnect each machine from the cluster before upgrading it. Then, other machines can continue working as a cluster until each is disconnected and upgraded.

**Procedure**

- Step 1** On a machine in the cluster, use the `disconnect` operation of `clusterconfig`. For example, to disconnect the machine `losangeles.example.com`, type `clusterconfig disconnect losangeles.example.com`. No `commit` is necessary.
- Step 2** Optionally, use the `suspendlistener` command to halt acceptance of new connections and messages during the upgrade process.
- Step 3** Issue the `upgrade` command to upgrade AsyncOS to a newer version.

**Note**

Disregard any warnings or confirmation prompts about disconnecting all of the machines in the cluster. Because you have disconnected the machine, AsyncOS does not disconnect the other machines in the cluster at this point.

- Step 4** Select the version of AsyncOS for the machine. The machine will reboot after the upgrade is complete.
- Step 5** Use the `resume` command on the upgraded machine to begin accepting new messages.
- Step 6** Repeat steps 1 - 5 for each machine in the cluster.

**Note**

After you disconnect a machine from the cluster, you cannot use it to change the configurations of other machines. Although you can still modify the cluster configuration, do not change it while machines are disconnected because settings can become unsynchronized.

- Step 7** After you have upgraded all the machines, use the `reconnect` operation of `clusterconfig` for each upgraded machine to reconnect it. For example, to reconnect the machine `losangeles.example.com`, type `clusterconfig reconnect losangeles.example.com`. Note that you can only connect a machine to a cluster that is running the same version of AsyncOS.

## CLI Command Support

### All Commands Are Cluster-aware

All CLI commands in AsyncOS are now cluster-aware. The behavior of some commands will change slightly when issued in a cluster mode. For example, the behavior of the following commands changes when issued on a machine that is part of a cluster:

## The commit and clearchanges Commands

### commit

The `commit` command commits all changes for all three levels of the cluster, regardless of which mode you are currently in.

### commitdetail

The `commitdetail` command provides details about configuration changes as they are propagated to all machines within a cluster.

### clearchanges

The `clearchanges` (`clear`) command clears all changes for all three levels of the cluster, regardless of which mode you are currently in.

## New Operation Added

### CLUSTERSHOW

Within each command, there is now a `CLUSTERSHOW` operation that allows you to see in which modes a command is configured.

When you enter a CLI command to perform an action that will be overridden by existing settings at a lower level, you will be presented with a notification. For example, if you are in cluster mode and enter a command, you may see a notification like this:

Note: Changes to these settings will not affect the following groups and machines because they are overriding the cluster-wide settings:

```
East_Coast, West_Coast
```

```
facilities_A, facilities_B, receiving_A
```

A similar message would be printed if you are editing settings for a group mode.

## Restricted Commands

Most CLI commands and their corresponding GUI pages can be run in any mode (cluster, group, or machine). However, some commands and pages are restricted to one mode only.

The system interface (either the GUI and the CLI) will always make it clear that a command is restricted and how it is restricted. It is easy to switch to the appropriate mode for configuring the command.

- In the GUI, use the “Change Mode” menu or the “Settings for this features are currently defined at:” links to switch modes.

- In the CLI, use the `clustermode` command to switch modes.

**Table 39-2** *Commands Restricted to Cluster Mode*

|                            |                         |
|----------------------------|-------------------------|
| <code>clusterconfig</code> | <code>sshconfig</code>  |
| <code>clustercheck</code>  | <code>userconfig</code> |
| <code>passwd</code>        |                         |

If you try to run one of these commands in group or machine mode, you will be given a warning message and the opportunity to switch to the appropriate mode.

**Note**

The `passwd` command is a special case because it needs to be usable by guest users. If a guest user issues the `passwd` command on a machine in a cluster, it will not print the warning message but will instead just silently operate on the cluster level data without changing the user's mode. All other users will get the above written behavior (consistent with the other restricted configuration commands).

The following commands are restricted to *machine mode*:

|                               |                               |                           |                              |
|-------------------------------|-------------------------------|---------------------------|------------------------------|
| <code>antispamstatus</code>   | <code>etherconfig</code>      | <code>resume</code>       | <code>suspenddel</code>      |
| <code>antispamupdate</code>   | <code>featurekey</code>       | <code>resumedel</code>    | <code>suspendlistener</code> |
| <code>antivirusstatus</code>  | <code>hostrate</code>         | <code>resumelister</code> | <code>techsupport</code>     |
| <code>antivirusupdate</code>  | <code>hoststatus</code>       | <code>rollovernow</code>  | <code>tophosts</code>        |
| <code>bouncerecipients</code> | <code>interfaceconfig</code>  | <code>routeconfig</code>  | <code>topin</code>           |
| <code>deleterecipients</code> | <code>ldapflush</code>        | <code>sbstatus</code>     | <code>trace</code>           |
| <code>delivernow</code>       | <code>ldaptest</code>         | <code>setgateway</code>   | <code>version</code>         |
| <code>diagnostic</code>       | <code>nslookup</code>         | <code>sethostname</code>  | <code>vofflush</code>        |
| <code>dnsflush</code>         | <code>quarantineconfig</code> | <code>settime</code>      | <code>vofstatus</code>       |
| <code>dnslistflush</code>     | <code>rate</code>             | <code>shutdown</code>     | <code>workqueue</code>       |
| <code>dnslisttest</code>      | <code>reboot</code>           | <code>status</code>       |                              |
| <code>dnsstatus</code>        | <code>resetcounters</code>    | <code>suspend</code>      |                              |

If you try to run one of the commands above in cluster or group mode, you will be given a warning message and the opportunity to switch to an appropriate mode.

The following commands are further restricted to the *login host* (i.e., the specific machine you are logged into). These commands require access to the local file system.

**Table 39-3** *Commands Restricted to Login Host Mode*

|                   |                             |                     |                      |
|-------------------|-----------------------------|---------------------|----------------------|
| <code>last</code> | <code>resetconfig</code>    | <code>tail</code>   | <code>upgrade</code> |
| <code>ping</code> | <code>supportrequest</code> | <code>telnet</code> | <code>who</code>     |

## Administering a Cluster from the GUI

Although you cannot create or join clusters or administer cluster specific settings from the GUI (the equivalent of the `clusterconfig` command), you can browse machines in the cluster, create, delete, copy, and move settings among the cluster, groups, and machines (that is, perform the equivalent of the `clustermode` and `clusterset` commands) from within the GUI.

The Incoming Mail Overview page is an example of a command that is restricted to the login host, because the Mail Flow Monitoring data you are viewing is stored on the local machine. To view the Incoming Mail Overview reports for another machine, you must log into the GUI for that machine.

Note the URL in the browser's address field when clustering has been enabled on an appliance. The URL will contain the word `machine`, `group`, or `cluster` as appropriate. For example, when you first log in, the URL of the Incoming Mail Overview page will appear as:

```
https://hostname/machine/serial_number/monitor/incoming_mail_overview
```

**Note**

The Incoming Mail Overview and Incoming Mail Details pages on the Monitor menu are restricted to the login machine.

The Mail Policies, Security Services, Network, and System Administration tabs contain pages that are not restricted to the local machine. If you click the Mail Policies tab, the centralized management information in the GUI changes.

**Figure 39-2 Centralized Management Feature in the GUI: No Settings Defined**

**Incoming Mail Policies** **Mode Indicator**

Mode — **Machine:example.com** Change Mode...

Centralized Management Options

Inheriting settings from Cluster: americas:

» Override Settings

Settings for this feature are currently defined at:

- Cluster: americas

**Find Policies**

Email Address:  Recipient Find Policies

☐ Sender

**Policies**

Add Policy...

| Order | Policy Name    | Anti-Spam                                            | Anti-Virus                                                                              | Virus Outbreak Filters | Content Filters | Delete |
|-------|----------------|------------------------------------------------------|-----------------------------------------------------------------------------------------|------------------------|-----------------|--------|
|       | Default Policy | IronPort<br>Positive: Deliver<br>Suspected: Disabled | Repaired: Deliver<br>Encrypted: Deliver<br>Unscannable: Deliver<br>Virus Positive: Drop | Enabled                | Disabled        |        |

Key: Default Custom Disabled

**Centralized Management box**


**Inherited settings (preview display)**

In [Figure 39-2](#), the machine is inheriting all of its configuration settings for the current feature from the cluster mode. The settings being inherited in a light grey (preview). You can retain these settings or change them, overriding the cluster level settings for this machine.

**Note**


The inherited settings (preview display) will always show the settings inherited from the cluster. Use caution when enabling or disabling dependent services among group and cluster levels. For more information, see [Copying and Moving Settings, page 39-11](#).

If you click the Override Settings link, you are taken to a new page for that feature. This page allows you to create new configuration settings for machine mode. You may begin with the default settings, or, if you've already configured settings in another mode, you can copy those settings to this machine.

**Figure 39-3** Centralized Management Feature in the GUI: Create New Settings


Alternatively, as shown in Figure 39-2, you can also navigate to modes where this configuration setting is already defined. The modes are listed in the lower half of the centralized management box, under “Settings for this feature are currently defined at:”. Only those modes where the settings are actually defined will be listed here. When you view a page for settings that are defined in (and inherited from) another mode, the page will display those settings for you.

If you click on one of the listed modes (for example, the Cluster: Americas link as shown in Figure 39-2), you will be taken to a new page that allows you to view and manage the settings for that mode.


**Figure 39-4** Centralized Management Feature in GUI: Settings Defined

When settings are defined for a given mode, the centralized management box is displayed on every page in a minimized state. Click the “Centralized Management Options” link to expand the box to show a list of options available for the current mode with respect to the current page. Clicking the “Manage Settings” button allows you to copy or move the current settings to a different mode or to delete those settings completely.

For example, in Figure 39-5, the Centralized Management Options link has been clicked to present the available options.

**Figure 39-5** Centralized Management Feature in GUI: Manage Settings

On the right side of the box is the “Change Mode” menu. This menu displays your current mode and provides the ability to navigate to any other mode (cluster, group, or machine) at any time.


**Figure 39-6** The Change Mode Menu  
Incoming Mail Policies

When you navigate to a page that represents a different mode, the “Mode —” text on the left side of the centralized management box will flash yellow, briefly, to alert you that your mode has changed.



Some pages within certain tabs are restricted to machine mode. However, unlike the Incoming Mail Overview page (which is restricted to the current login host), these pages can be used for any machine in the cluster.

**Figure 39-7 Centralized Management Feature: Machine Restricted**



Choose which machine to administer from the Change Mode menu. You will see a brief flashing of the text to remind you that you have changed modes.

## Cluster Communication

Machines within a cluster communicate with each other using a *mesh network*. By default, all machines connect to all other machines. If one link goes down, other machines will not be prevented from receiving updates.

By default, all intra-cluster communication is secured with SSH. Each machine keeps an in-memory copy of the route table and makes in-memory changes as necessary if links go down or up. Each machine also performs a periodic “ping” (every 1 minute) of every other machine in the cluster. This ensures up-to-date link status and maintains the connections in case a router or NAT has a timeout.



### Note

The connection between two clustered appliances may be dropped if one of the appliances attempts to open more than the maximum number of SSH connections allowed. The appliances automatically rejoin the cluster within seconds and no manual configuration is needed.

## DNS and Hostname Resolution

DNS is required to connect a machine to the cluster. Cluster communication is normally initiated using the DNS hostnames of the machines (not the hostname of an interface on the machine). A machine with an unresolvable hostname would be unable to actually communicate with any other machines in the cluster, even though it is technically part of the cluster.

Your DNS must be configured to have the hostname point to the correct IP interface on the appliance that has SSH or CCS enabled. This is very important. If DNS points to another IP address that does not have SSH or CCS enabled it will not find the host. Note that centralized management uses the “main hostname,” as set with the `sethostname` command, not the per-interface hostname.

If you use an IP address to connect to another machine in the cluster, the machine you connect to must be able to make a reverse look up of the connecting IP address. If the reverse look up times out because the IP address isn’t in the DNS, the machine cannot connect to the cluster.

## Clustering, Fully Qualified Domain Names, and Upgrading

DNS changes can cause a loss of connectivity after upgrading AsyncOS. Please note that if you need to change the fully qualified domain name of a machine in the cluster (not the hostname of an interface on a machine in the cluster), you must change the hostname settings via `sethostname` and update the DNS record for that machine *prior* to upgrading AsyncOS.

## Cluster Communication Security

Cluster Communication Security (CCS) is a secure shell service similar to a regular SSH service. Cisco implemented CCS in response to concerns regarding using regular SSH for cluster communication. SSH communication between two machines opens regular logins (admin, etc.) on the same port. Many administrators prefer not to open regular logins on their clustered machines.

Tip: never enable Cluster Communication Services, even though it is the default, unless you have firewalls blocking port 22 between some of your clustered machines. Clustering uses a full mesh of SSH tunnels (on port 22) between all machines. If you have already answered Yes to enabling CCS on any machine, remove all machines from the cluster and start again. Removing the last machine in the cluster removes the cluster.

CCS provides an enhancement where the administrator can open up cluster communication, but not CLI logins. By default, the service is disabled. You will be prompted to enable CCS from the `interfaceconfig` command when you are prompted to enable other services. For example:

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

```
Which port do you want to use for SSH?
```

```
[22]>
```

```
Do you want to enable Cluster Communication Service on this interface?
```

```
[N]> y
```

```
Which port do you want to use for Cluster Communication Service?
```

```
[2222]>
```

The default port number for CCS is 2222. You may change this to another open, unused, port number if you prefer. After the join is complete and the joining machine has all the configuration data from the cluster, the following question is presented:

```
Do you want to enable Cluster Communication Service on this interface? [N]> y
```

```
Which port do you want to use for Cluster Communication Service?
```

```
[2222]>
```

## Cluster Consistency

The machines that are “cluster aware” will continually verify network connections to other machines within the cluster. This verification is done by periodic “pings” sent to other machines in the cluster.

If all attempts to communicate with a particular machine fail, then the machine that has been trying to communicate will log a message saying that the remote host has disconnected. The system will send an alert to the administrator that the remote host went down.

Even if a machine is down, the verification pings will continue to be sent. When a machine rejoins the cluster network, a synchronization command will be issued so that any previously offline machines can download any updates. The synchronization command will also determine if there have been any changes on one side but not the other. If so, then the previously down machine will silently download the updates.

## Disconnect/Reconnect

A machine may be disconnected from a cluster. Occasionally, you may intend to deliberately disconnect the machine, for example, because you are upgrading the machine. A disconnect could also occur by accident, for example, due to a power failure or other software or hardware error. A disconnect can also occur if one appliance attempts to open more than the maximum number of SSH connections allowed in a session. A machine that is disconnected from a cluster can still be accessed directly and configured; however, any changes made will not be propagated to other machines within the cluster until the disconnected machine becomes reconnected.

When a machine reconnects to the cluster, it tries to reconnect to all machines at once.

In theory, two machines in a cluster that are disconnected could commit a similar change to their local databases at the same time. When the machines are reconnected to the cluster, an attempt will be made to synchronize these changes. If there is a conflict, the most recent change is recorded (supersedes any other changes).

During a commit, the appliance checks every variable that is being changed. The commit data includes version information, sequence identification numbers, and other information that can be compared. If the data you are about to change is found to be in conflict with previous changes, you will be given the option to discard your changes. For example, you might see something like this:

```
(Machine mail3.example.com)> clustercheck
```

```
This command is restricted to "cluster" mode. Would you like to switch to "cluster" mode? [Y]> y
```

```
Checking Listeners (including HAT, RAT, bounce profiles)...
```

```
Inconsistency found!
```

```
Listeners (including HAT, RAT, bounce profiles) at Cluster enterprise:
```

```
mail3.example.com was updated Mon Sep 12 10:59:17 2005 PDT by 'admin' on mail3.example.com
```

```
test.example.com was updated Mon Sep 12 10:59:17 2005 PDT by 'admin' on mail3.example.com
```

```
How do you want to resolve this inconsistency?
```

1. Force entire cluster to use test.example.com version.
  2. Force entire cluster to use mail3.example.com version.
  3. Ignore.
- [1]>

If you choose not to discard your changes, they are still intact (but uncommitted). You can review your changes against the current settings and decide how to proceed.

You can also use the `clustercheck` command at any time to verify that the cluster is operating correctly.

```
losangeles> clustercheck
```

```
Do you want to check the config consistency across all machines in the cluster? [Y]> y
```

```
Checking losangeles...
```

```
Checking newyork...
```

```
No inconsistencies found.
```

## Interdependent Settings

In a centrally managed environment, some interdependent settings are configured in different modes. The flexibility of the configuration model allows you to configure settings at multiple modes, and the laws of inheritance govern which settings will be used on a per-machine basis. However, some settings have dependencies on other settings, and the availability of the dependent settings' configuration is not limited to settings at the same mode. Thus, it is possible to configure a setting for one level that references a setting that is configured for a specific machine at a different level.

The most common example of an interdependent setting involves a select field on a page that pulls data from a different cluster section. For example, the following features can be configured in different modes:

- using LDAP queries
- using dictionaries or text resources
- using bounce or SMTP authentication profiles.

Within centralized management, there are restricted and non-restricted commands. (See [Restricted Commands, page 39-14.](#)) Non-restricted commands are generally configuration commands that can be shared across the cluster.

The `listenerconfig` command is an example of a command that can be configured for all machines in a cluster. Non-restricted commands represent commands that can be mirrored on all machines in a cluster, and do not require machine-specific data to be modified.

Restricted commands, on the other hand, are commands that only apply to a specific mode. For example, users cannot be configured for specific machines — there must be only one user set across the whole cluster. (Otherwise, it would be impossible to login to remote machines with the same login.) Likewise, since the Mail Flow Monitor data, System Overview counters, and log files are only maintained on a per-machine basis, these commands and pages must be restricted to a machine.

You will notice that while Scheduled Reports may be configured identically across the whole cluster, the viewing of reports is machine-specific. Therefore, within a single Scheduled Reports page in the GUI, configuration must be performed at the cluster mode, but viewing of reports must be done at the machine mode.

The System Time pages encompass the `settz`, `ntpconfig`, and `settime` commands, and thus represents a mixture of restricted and non-restricted commands. In this case, `settime` must be restricted to machine-only modes (since time settings are specific for machine), while `settz` and `ntpconfig` may be configured at cluster or group modes.

**Figure 39-8 Example of Interdependent Settings**

#### Edit Listener

Mode — Cluster: americas Change Mode

Centralized Management Options

**Listener Settings**

|                               |                                                                                              |
|-------------------------------|----------------------------------------------------------------------------------------------|
| Name:                         | IncomingMail                                                                                 |
| Type of Listener:             | Public                                                                                       |
| Interface:                    | Data 1 TCP Port: 25                                                                          |
| Bounce Profile:               | Default                                                                                      |
| Disclaimer Above:             | None<br><small>Disclaimer text will be applied above the message body.</small>               |
| Disclaimer Below:             | <div> None<br/> None<br/> <b>disclaimer (- Unavailable on Machine: buttercup.run)</b> </div> |
| SMTP Authentication Profile:  |                                                                                              |
| Certificate:                  | test                                                                                         |
| SMTP Address Parsing Options: | Optional settings for controlling parsing in SMTP "MAIL FROM" and "RCPT"                     |
| Advanced:                     | Optional settings for customizing the behavior of the Listener                               |
| LDAP Queries:                 | Optional settings for controlling LDAP queries associated with this Listener                 |

In this representation, the listener “IncomingMail” is referencing a footer named “disclaimer” that has been configured at the machine level only. The drop-down list of available footer resources shows that the footer is not available on the machine “buttercup.run” which is also available in the cluster. There are two solutions to this dilemma:

- promote the footer “disclaimer” from the machine level to the cluster level
- demote the listener to the machine level to remove the interdependency

In order to fully maximize the features of a centrally managed system, the former solution is preferred. Be aware of interdependencies among settings as you tailor the configuration of your clustered machines.

## Loading a Configuration in Clustered Appliances

AsyncOS for Email allows you to load a cluster configuration in clustered appliances. You can load the cluster configuration in the following scenarios:

- If you are migrating from an on-premise environment to a hosted environment and you want to migrate the on-premise cluster configuration to the hosted environment.

- If an appliance in a cluster is down or needs to be retired and you want to load the configuration from this appliance to a new appliance that you plan to add to the cluster.
- If you are adding more appliances to your cluster and you want to load the configuration from one of the existing appliances in the cluster to the newly added appliances.
- If you want to load a backed-up configuration to a cluster.

Depending on your requirements, you can load a cluster configuration or appliance configuration from a valid cluster configuration file.

**Note**

You cannot load the configuration of a standalone appliance on a clustered appliance.

**Before You Begin**

- Make sure that you have a valid and complete XML configuration. See [Loading a Configuration File, page 33-9](#).
- Create a backup of the current configuration of the appliance to which you plan to load the configuration. See [Saving and Exporting the Current Configuration File, page 33-8](#).
- Create a cluster setup with all the appliances that you plan to have in your setup. See [Creating and Joining a Cluster, page 39-4](#).

**Note**

You can have all the appliances under one group. Ensure that the interfaces for cluster communication in your setup have same names and SSH and CCS settings as in the XML configuration.

**Procedure**

- 
- Step 1** Click **System Administration > Configuration File**.
- Step 2** Choose the cluster from the **Mode** drop-down menu.
- Step 3** Depending on whether you want to load cluster or appliance configuration, do one of the following:
- **Load Cluster Configuration**
    - a. In the Load Configuration section, choose **Cluster** from the drop-down list.
    - b. Load the cluster configuration, and click **Load**. See [Loading a Configuration File, page 33-9](#).
    - c. Assign groups from the loaded configuration to the appliances in the cluster, and copy appliance configuration from the appliances in the selected group to the respective appliances. Use the **Group Configuration** and **Appliance Configuration** drop-down lists.

If you do not want to copy an appliance configuration, choose **Don't Copy** from the **Appliance Configuration** drop-down list.
    - d. Review the configuration. Click **Review**.
    - e. Click **Confirm**.
    - f. Click **Continue**.
  - **Load Appliance Configuration**
    - a. In the Load Configuration section, choose **Appliance in cluster** from the drop-down list.
    - b. Load the configuration, and click **Load**. See [Loading a Configuration File, page 33-9](#). Note that you cannot load the configuration of a standalone appliance on a clustered appliance.

- c. Choose the appliance configuration from the loaded configuration and the intended appliance in the cluster to which you want to load the configuration. Use the drop-down lists.
- d. Click **OK**.
- e. Click **Continue**.
- f. To load the appliance configuration to more appliances, repeat Step a through Step e.

**Step 4** Review the network settings of the clustered appliances, and commit your changes.

---

## Best Practices and Frequently Asked Questions

### Best Practices

When you create the cluster, the machine you happen to be logged into is automatically added to the cluster as the first machine, and also added to the Main\_Group. Its machine level settings effectively get moved to the cluster level as much as possible. There are no settings at the group level, and the only settings left at the machine level are those which do not make sense at the cluster level, and cannot be clustered. Examples are IP addresses, featurekeys, etc.

Leave as many settings at the cluster level as possible. If only one machine in the cluster needs a different setting, copy that cluster setting to the machine level for that machine. Do not move that setting. If you move a setting which has no factory default (e.g. HAT table, SMTPROUTES table, LDAP server profile, etc.), the systems inheriting the cluster settings will have blank tables and will probably not process email.

To have that machine re-inherit the cluster setting, manage the CM settings and delete the machine setting. You will only know if a machine is overriding the cluster setting when you see this display:

Settings are defined:

To inherit settings from a higher level: Delete Settings for this feature at this mode.

You can also Manage Settings.

Settings for this feature are also defined at:

Cluster: xxx

Or this display:

Delete settings from:

Cluster: xxx

Machine: yyyy.domain.com

### Copy vs. Move

When to copy: when you want the cluster to have a setting, and a group or machine to also have no settings or to have different settings.

When to move: when you want the cluster to have no setting at all, and for the group or machine to have the settings.

## Good CM Design Practices

When you LIST your CM machines, you want to see something like this:

```
cluster = CompanyName

Group Main_Group:

 Machine lab1.example.com (Serial #: XXXXXXXXXXXX-XXXXXXX)
 Machine lab2.example.com (Serial #: XXXXXXXXXXXX-XXXXXXX)

Group Paris:

 Machine lab3.example.com (Serial #: XXXXXXXXXXXX-XXXXXXX)
 Machine lab4.example.com (Serial #: XXXXXXXXXXXX-XXXXXXX)

Group Rome:

 Machine lab5.example.com (Serial #: XXXXXXXXXXXX-XXXXXXX)
 Machine lab6.example.com (Serial #: XXXXXXXXXXXX-XXXXXXX)
```

Be careful not to lose track of the level at which you are making changes. For example, if you have changed the name of your Main\_Group (using RENAMEGROUP) to London, it will look like this:

```
cluster = CompanyName

Group London:

 Machine lab1.cable.nu (Serial #: 000F1FF7B3F0-CF2SX51)
 ...
```

However, this configuration tends to confuse many administrators, because they begin making changes to the London systems at the group level, and they stop using the Cluster level as the normal configuration level for basic settings.

**Tip:** it is not a good practice to have a group with the same name as the cluster, e.g. cluster London, group London. If you are using site names for group names, it is not good practice to have a cluster name that refers to a location.

The correct method, as explained above, is to leave as many settings at the cluster level as possible. In most cases you should leave your primary site or main collection of machines in the Main\_Group, and use groups for your additional sites. This is true even if you consider that both sites are “equal.” Remember, CM has no primary/secondary or master/slave servers — all clustered machines are peers.

**Tip:** if you will be using extra groups you can easily prepare the groups before those extra machines are joined to the cluster.

## Best Practices for Accessing Spam Quarantines in Cluster Setup

Accessing spam quarantines of other appliances in a cluster from the logged-in appliance may cause excessive CPU usage on the logged-in appliance. To avoid this scenario, you can access the spam quarantines by logging into the respective appliances.

## Procedures: Configuring an Example Cluster

To configure this example cluster, log out of all GUIs on all machines before running `clusterconfig`. Run `clusterconfig` on any one of the primary site machines. You will then join to this cluster only the other local and remote machines that need the maximum possible shared settings (allowing for the



machine only-settings like IP address). The `clusterconfig` command cannot be used to join a remote machine to the cluster — you must use the CLI on the remote machine and run `clusterconfig` (“join an existing cluster”).

In our example above we log in to lab1, run `clusterconfig` and create a cluster called CompanyName. We have only one machine with identical requirements, so we log in to lab2, and `saveconfig` the existing configuration (it will be drastically altered when it inherits most of lab1 settings.) On lab2 we can then use `clusterconfig` to join an existing cluster. Repeat if you have additional machines at this site needing similar policies and settings.

Run `CONNSTATUS` to confirm that DNS resolves correctly. As machines are joined to the cluster, the new machines inherit almost all of their settings from lab1 and their older settings are lost. If they are production machines you will need to anticipate if mail will still be processed using the new configuration instead of their previous configuration. If you remove them from the cluster, they will not revert to their old, private configs.

Next, we count the number of exceptional machines. If there is only one, it should receive a few extra machine level settings and you will not need to create an extra group for it. Join it to the cluster and begin copying settings down to the machine level. If this machine is an existing production machine you must back up the configuration and consider the changes to mail processing as above.

If there are two or more, as in our example, decide if those two will share any settings with each other that are not shared with the cluster. In that case, you will be creating one or more groups for them. Otherwise, you will make machine level settings for each, and do not need to have extra groups.

In our case we want to run `clusterconfig` from the CLI on any of the machines already in the cluster, and select `ADDGROUP`. We will do this twice, once for Paris and once for Rome.

Now you can begin using the GUI and CLI to build configuration settings for the cluster and for ALL the groups, even if the groups have no machines in them yet. You will only be able to create machine specific settings for machines *after* they have joined the cluster.

The best way to create your override or exceptional settings is to copy the settings from the higher (e.g. cluster) level down to a lower (e.g. group) level.

For example, after creating the cluster our `dnsconfig` settings initially looked like this:

```
Configured at mode:
Cluster: Yes
Group Main_Group: No
Group Paris: No
Group Rome: No
Machine lab2.cable.nu: No
```

If we "Copy to Group" the DNS settings, it will look like this:

```
Configured at mode:
Cluster: Yes
Group Main_Group: No
Group Paris: Yes
Group Rome: No
Machine lab2.cable.nu: No
```

Now you can edit the Paris group-level DNS settings, and other machines in the Paris group will inherit them. Non-Paris machines will inherit the cluster settings, unless they have machine-specific settings. Besides DNS settings, it is common to create group level settings for SMTPROUTES.

Tip: when using the CLI CLUSTERSET function in various menus, you can use a special option to copy settings to All Groups, which is not available through the GUI.

Tip: complete listeners will be automatically inherited from the group or cluster, and you normally only create these on the first system in the cluster. This reduces administration considerably. However, for this to work *you must name the Interfaces identically throughout your group or cluster*.

Once the settings are defined correctly at the group level, you can join machines to the cluster and make them part of this group. This requires two steps:

First, to join our remaining 4 systems to the cluster, we run `clusterconfig` on each. The larger and more complex the cluster, the longer it takes to join, and this can take several minutes. You can monitor the joining progress with the LIST and CONNSTATUS sub-commands. After the joins are complete you can use SETGROUP to move the machines from the Main\_Group into Paris and Rome. There is no way to avoid the fact that initially, all machines added to the cluster inherit the Main\_Group settings, not the Paris and Rome settings. This could affect mail flow traffic if the new systems are already in production.

Tip: do not make your lab machines part of the same cluster as your production machines. Use a new cluster name for lab systems. This provides an added layer of protection against unexpected changes (someone changing a lab system and accidentally losing production mail, for example).

## Summary of GUI Options for Using CM Settings Other Than the Cluster Default

Override settings, and start with default settings. For example, the default settings for the SMTPROUTES configuration is a blank table, which you can then build from scratch.

Override settings, but start with a copy of the settings currently inherited from Cluster xxx, or group yyy. For example, you may want a new copy of the SMTPROUTES table at the group level which is initially identical to the cluster table. All Cisco appliances that are contained in that same group (SETGROUP) will get this table. Machines not in the group will still use the cluster level settings. Changing the SMTPROUTES on this independent copy of the table will not affect other groups, machines inheriting the cluster settings, or machines where the setting is defined at the individual machine level. This is the most common selection.

Manage settings, a sub-menu of Centralized Management Options. From this menu you can copy as above, but you can also move or delete settings. If you move the SMTPROUTES to a group or machine level, then the routes table will be blank at the cluster level but will exist at the more specific level.

Manage settings. Continuing our SMTPROUTES example, using the delete option will also result in a blank SMTPROUTES table for the cluster. This is fine if you previously configured definitions for SMTPROUTES at the group level or machine levels. It is not a best practice to delete the cluster level settings and rely only on group or machine settings. The cluster-wide settings are useful as defaults on newly added machines, and keeping them reduces the number of group or site settings you have to maintain by one.

## Setup and Configuration Questions

Q. I have a previously configured standalone machine and I join an existing cluster. What happens to my settings?

A. When a machine joins a cluster, all of that machine's clusterable settings will be inherited from the cluster level. Upon joining a cluster, all locally configured non-network settings will be lost, overwritten with the settings of the cluster and any associated groups. (This includes the user/password table; passwords and users are shared within a cluster).

Q. I have a clustered machine and I remove it (permanently) from the cluster. What happens to my settings?

A. When a machine is permanently removed from a cluster, its configuration hierarchy is “flattened” such that the machine will continue to work the same as it did when it was part of the cluster. All settings that the machine has been inheriting will be applied to the machine in the standalone setting.

For example, if there is only a cluster-mode Global Unsubscribe table, that Global Unsubscribe table data will be copied to the machine's local configuration when the machine is removed from the cluster.

## General Questions

Q. Are log files aggregated within centrally managed machines?

A. No. Log files are still retained for each individual machines. The Security Management appliance can be used to aggregate mail logs from multiple machines for the purposes of tracking and reporting.

Q. How does User Access work?

A. The Cisco appliances share one database for the entire cluster. In particular, there is only `admin` account (and password) for the entire cluster.

Q. How should I cluster a data center?

A. Ideally, a data center would be a “group” within a cluster, not its own cluster. However, if the data centers do not share much between themselves, you may have better results with separate clusters for each data center.

Q. What happens if systems are offline and they reconnect?

A. Systems attempt to synchronize upon reconnecting to the cluster.

## Network Questions

Q. Is the centralized management feature a “peer-to-peer” architecture or a “master/slave” architecture?

A. Because every machine has all of the data for all of the machines (including all machine-specific settings that it will never use), the centralized management feature can be considered a peer-to-peer architecture.

Q. How do I set up a box so it is not a peer? I want a “slave” system.

A. Creating a true “slave” machine is not possible with this architecture. However, you can disable the HTTP (GUI) and SSH/Telnet (CLI) access at the machine level. In this manner, a machine without GUI or CLI access *only* be configured by `clusterconfig` commands (that is, it can never be a login host). This is similar to having a slave, but the configuration can be defeated by turning on login access again.

Q. Can I create multiple, segmented clusters?

A. Isolated “islands” of clusters are possible; in fact, there may be situations where creating them may be beneficial, for example, for performance reasons.

Q. I would like to reconfigure the IP address and hostname on one of my clustered appliances. If I do this, will I lose my GUI/CLI session before being able to run the reboot command?

Follow these steps:

- a. Add the new IP address
- b. Move the listener onto the new address
- c. Leave the cluster
- d. Change the hostname
- e. Make sure that oldmachinename does not appear in the `clusterconfig` connections list when viewed from any machine
- f. Make sure that all GUI sessions are logged out
- g. Make sure that CCS is not enabled on any interface (check via `interfaceconfig` or Network > Listeners)
- h. Add the machine back into the cluster

Q. Can the Destination Controls function be applied at the cluster level, or is it local machine level only?

It may be set at a cluster level; however, the limits are on a per-machine basis. So if you limit to 50 connections, that is the limit set for each machine in the cluster.

## Planning and Configuration

Q. What can I do to maximize efficiency and minimize problems when setting up a cluster?

1. Initial Planning

- Try to configure as many things as possible at the cluster level.
- Manage by machines only for the exceptions.
- If you have multiple data centers, for example, use groups to share traits that are neither cluster-wide nor necessarily machine-specific.
- Use the same name for Interfaces and Listeners on each of the appliances.

2. Be aware of restricted commands.

3. Pay attention to interdependencies among settings.

For example, the `listenerconfig` command (even at the cluster level) depends on interfaces that only exist at a machine level. If the interface does not exist at the machine level on all machines in the cluster, that listener will be disabled.

Note that deleting an interface would also affect `listenerconfig`.

4. Pay attention to your settings!

Remember that previously-configured machines will lose their independent settings upon joining a cluster. If you want to re-apply some of these previously configured settings at the machine level, be sure to take note of all settings before joining the cluster.

Remember that a “disconnected” machine is still part of the cluster. When it is reconnected, any changes you made while it was offline will be synchronized with the rest of the cluster.

Remember that if you permanently remove a machine from a cluster, it will retain all of the settings it had as part of that cluster. However, if you change your mind and re-join the cluster, the machine will lose all standalone settings.

Use the `saveconfig` command to keep records of settings.



## Testing and Troubleshooting

---

- [Debugging Mail Flow Using Test Messages: Trace, page 40-1](#)
- [Using the Listener to Test the Appliance, page 40-12](#)
- [Troubleshooting the Network, page 40-16](#)
- [Troubleshooting the Listener, page 40-22](#)
- [Troubleshooting Email Delivery From the Appliance, page 40-23](#)
- [Troubleshooting Performance, page 40-26](#)
- [Responding to Alerts, page 40-27](#)
- [Remotely Resetting Appliance Power, page 40-27](#)
- [Working with Technical Support, page 40-28](#)



### Note

Several of the features or commands described in this section will affect, or be affected by routing precedence. Please see [Assigning Network and IP Addresses](#) for more information.

---

## Debugging Mail Flow Using Test Messages: Trace

You can use System Administration > Trace page (the equivalent of the `trace` command in the CLI) to debug the flow of messages through the system by emulating sending a test message. The Trace page (and `trace` CLI command) 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 of the system (*including uncommitted changes*). The test message is not actually sent. The Trace page (and `trace` CLI command) can be a powerful troubleshooting or debugging tool, especially if you have combined many of the advanced features available on the Cisco appliance.



### Note

Trace is not effective for testing file reputation scanning.

---

The Trace page (and `trace` CLI command) prompts you for the input parameters listed in [Table 40-1](#).

**Table 40-1** *Input for the Trace page*

| Value                                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Example                                                                                                                                                       |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Source IP address                            | Type the IP address of the remote client to mimic the source of the remote domain. This can be an Internet Protocol version 4 (IPv4) or version 6 (IPv6) address.<br><br><b>Note:</b> The <code>trace</code> command prompts for an IP address and a fully-qualified domain name. It does <i>not</i> attempt to reverse the IP address to see if it matches the fully-qualified domain name. The <code>trace</code> command does not allow the fully-qualified domain name field to be blank, so it is impossible to test a scenario where the DNS does not reverse match properly. | <code>203.45.98.109</code><br><br><code>2001:0db8:85a3::8a2e:0370:7334</code>                                                                                 |
| Fully Qualified Domain Name of the Source IP | Type the fully-qualified remote domain name to mimic. If left null, a reverse DNS lookup will be performed on the source IP address.                                                                                                                                                                                                                                                                                                                                                                                                                                                | <code>smtp.example.com</code>                                                                                                                                 |
| Listener to Trace Behavior on                | Choose from the list of listeners configured on the system to emulate sending the test message to.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <code>InboundMail</code>                                                                                                                                      |
| SenderBase Network Owner Organization ID     | Type the unique identification number of the SenderBase network owner, or allow the system to Lookup network owner ID associated with source IP address.<br>You can view this information if you added network owners to sender groups via the GUI.                                                                                                                                                                                                                                                                                                                                 | <code>34</code>                                                                                                                                               |
| SenderBase Reputation Score (SBRS scores)    | Type the SBRS score you want to provide for the spoofed domain, or allow the system to look up the SBRS score associated with the source IP address. This can be helpful when testing policies that use SBRS scores. Note that manually entered SBRS scores are not passed to the Context Adaptive Scanning Engine (CASE). See <a href="#">Editing Sender Reputation Filtering Score Thresholds for a Listener</a> , page 6-5 for more information.                                                                                                                                 | <code>-7.5</code>                                                                                                                                             |
| Envelope Sender                              | Type the Envelope Sender of the test message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <code>admin@example.net</code>                                                                                                                                |
| Envelope Recipients                          | Type a list of recipients for the test message. Separate multiple entries with commas.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <code>joe</code><br><code>frank@example.com</code>                                                                                                            |
| Message Body                                 | Type the message body for the test message, including headers. Type a period on a separate line to end entering the message body. Note that “headers” are considered part of a message body (separated by a blank line), and omitting headers, or including poorly formatted ones can cause unexpected trace results.                                                                                                                                                                                                                                                               | <code>To: 1@example.com</code><br><br><code>From: ralph</code><br><code>Subject: Test</code><br><br><code>this is a test message</code><br><br><code>.</code> |

After you have entered the values, click **Start Trace**. A summary of all features configured on the system affecting the message is printed.

You can upload message bodies from your local file system. (In the CLI, you can test with message bodies you have uploaded to the `/configuration` directory. See [FTP, SSH, SCP, and Telnet Access](#) for more information on placing files for import onto the Cisco appliance.)

After the summary is printed, you are prompted to view the resulting message and re-run the test message again. If you enter another test message, the Trace page and the `trace` command uses any previous values from [Table 40-1](#) you entered.

**Note**

The sections of configuration tested by the `trace` command listed in [Table 40-2](#) are performed *in order*. This can be extremely helpful in understanding how the configuration of one feature affects another. For example, a recipient address transformed by the domain map feature will affect the address as it is evaluated by the RAT. A recipient that is affected by the RAT will affect the address as it is evaluated by alias table, and so on.

**Table 40-2**      **Viewing Output When Performing a Trace**

| trace Command Section                                                                                                                                                                                                                                                                | Output                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Host Access Table (HAT) and Mail Flow Policy Processing                                                                                                                                                                                                                              | <p>The Host Access Table settings for the listener you specified are processed. The system reports which entry in the HAT matched from the remote IP address and remote domain name you entered. You can see the default mail flow policies and sender groups and which one matched the given entries.</p> <p>If the Cisco appliance was configured to reject the connection (either through a REJECT or TCPREFUSE access rule), the <code>trace</code> command exits at the point in the processing.</p> <p>For more information on setting HAT parameters, see <a href="#">Understanding Predefined Sender Groups and Mail Flow Policies</a>, page 7-11.</p> |
| <b>Envelope Sender Address Processing</b>                                                                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| These sections summarize how the appliance configuration affects the Envelope Sender you supply. (That is, how the MAIL FROM command would be interpreted by the configuration of the appliance.) The <code>trace</code> command prints “Processing MAIL FROM:” before this section. |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Default Domain                                                                                                                                                                                                                                                                       | <p>If you specified that a listener to change the default sender domain of messages it receives, any change to the Envelope Sender is printed in this section.</p> <p>For more information, see <a href="#">Configuring the Gateway to Receive Email</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                  |



**Table 40-2** Viewing Output When Performing a Trace (continued)

| trace Command Section                                                                                                                                                                                                                                                                                                     | Output                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Masquerading                                                                                                                                                                                                                                                                                                              | <p>If you specified that the Envelope Sender of a message should be transformed, the change is noted here. You enable masquerading for the Envelope Sender on private listeners using the <code>listenerconfig -&gt; edit -&gt; masquerade -&gt; config</code> subcommands.</p> <p>For more information, see <a href="#">Configuring Routing and Delivery Features</a>.</p> |
| <b>Envelope Recipient Processing</b> <p>These sections summarize how the appliance affects the Envelope Recipients you supply. (That is, how the RCPT TO command would be interpreted by the configuration of the appliance.) The <code>trace</code> command prints “Processing Recipient List:” before this section.</p> |                                                                                                                                                                                                                                                                                                                                                                             |
| Default Domain                                                                                                                                                                                                                                                                                                            | <p>If you specified that a listener to change the default sender domain of messages it receives, any changes to the Envelope Recipients are printed in this section.</p> <p>For more information, see <a href="#">Configuring the Gateway to Receive Email</a>.</p>                                                                                                         |
| Domain Map Translation                                                                                                                                                                                                                                                                                                    | <p>The domain map feature transforms the recipient address to an alternate address. If you specified any domain map changes and a recipient address you specified matches, the transformation is printed in this section.</p> <p>For more information, see <a href="#">Configuring Routing and Delivery Features</a>.</p>                                                   |
| Recipient Access Table (RAT)                                                                                                                                                                                                                                                                                              | <p>Each Envelope Recipient that matches an entry in the RAT is printed in this section, in addition to the policy and parameters. (For example, if a recipient was specified to bypass limits in the listener’s RAT.)</p> <p>For more information on specifying recipients you accept, see <a href="#">Configuring the Gateway to Receive Email</a>.</p>                    |
| Alias Table                                                                                                                                                                                                                                                                                                               | <p>Each Envelope Recipient that matches an entry in the alias tables configured on the appliance (and the subsequent transformation to one or more recipient addresses) is printed in this section.</p> <p>For more information, see <a href="#">Configuring Routing and Delivery Features</a>.</p>                                                                         |

**Pre-Queue Message Operations**

These sections summarize how the appliance affects each message after the message contents have been received, but before the messages are enqueued on the work queue. This processing occurs before the final `250 ok` command is returned to the remote MTA.

The `trace` command prints “Message Processing:” before this section.

**Table 40-2**      **Viewing Output When Performing a Trace (continued)**

| trace Command Section | Output                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Virtual Gateways      | <p>The <code>altsrchost</code> command assigns messages to a specific interface, based on a match of the Envelope Sender's full address, domain, or name, or IP address. If an Envelope Sender matches entries from the <code>altsrchost</code> command, that information is printed in this section.</p> <p>Note that the virtual gateway address assigned at this point may be overridden by message filter processing below.</p> <p>For more information, see <a href="#">Configuring Routing and Delivery Features</a>.</p> |
| Bounce Profiles       | <p>Bounce profiles are applied at three different points in the processing. This is the first occurrence. If a listener has a bounce profile assigned to it, it is assigned at this point in the process. That information is printed in this section.</p> <p>For more information, see <a href="#">Configuring Routing and Delivery Features</a>.</p>                                                                                                                                                                          |

**Table 40-2 Viewing Output When Performing a Trace (continued)**

| trace Command Section                                                                                                                                                                                                                                                                                                                                                                                                          | Output                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Work Queue Operations</b><br><br>The following group of functions are performed on messages in the work queue. This occurs after the message has been accepted from the client, but before the message is enqueued for delivery on a destination queue. “Messages in Work Queue” is reported by the <code>status</code> and <code>status detail</code> commands.                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Masquerading                                                                                                                                                                                                                                                                                                                                                                                                                   | If you specified that the To:, From:, and CC: headers of messages should be masked (either from a static table entered from a listener or via an LDAP query), the change is noted here. You enable masquerading for the message headers on private listeners using the <code>listenerconfig -&gt; edit -&gt; masquerade -&gt; config</code> subcommands.<br><br>For more information, see <a href="#">Configuring Routing and Delivery Features</a> .                                                                                                                                                                                                  |
| LDAP Routing                                                                                                                                                                                                                                                                                                                                                                                                                   | If LDAP queries have been enabled on a listener, the results of LDAP acceptance, re-routing, masquerading, and group queries are printed in this section.<br><br>For more information, see <a href="#">LDAP Queries</a> .                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Message Filters Processing                                                                                                                                                                                                                                                                                                                                                                                                     | All messages filters that are enabled on the system are evaluated by the test message at this point. For each filter, the rule is evaluated, and if the end result is “true,” each of the actions in that filter are then performed in sequence. A filter may contain other filters as an action, and the nesting of filters is unlimited. If a rule evaluates to “false” and a list of actions is associated with an <code>else</code> clause, those actions are evaluated instead. The results of the message filters, processed in order, are printed in this section.<br><br>See <a href="#">Using Message Filters to Enforce Email Policies</a> . |
| <b>Mail Policy Processing</b><br><br>The mail policy processing section displays the Anti-Spam, Anti-Virus, Outbreak Filters feature, and disclaimer stamping for all recipients you supplied. If multiple recipients match multiple policies in Email Security Manager, the following sections will be repeated for each matching policy. The string: “Message Going to” will define which recipients matched which policies. |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

**Table 40-2** Viewing Output When Performing a Trace (continued)

| trace Command Section       | Output                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Anti-Spam                   | <p>This section notes messages that are not flagged to be processed by anti-spam scanning. If messages are to be processed by anti-spam scanning for the listener, the message is processed and the verdict returned is printed. If the Cisco appliance is configured to bounce or drop the messages based on the verdict, that information is printed and the <code>trace</code> command processing stops.</p> <p>Note: This step is skipped if anti-spam scanning is unavailable on the system. If anti-spam scanning is available but has not been enabled with a feature key, that information is also printed in this section.</p> <p>See <a href="#">Anti-Spam</a>.</p>                                                                                 |
| Anti-Virus                  | <p>This section notes messages that are not flagged to be processed by anti-virus scanning. If messages are to be processed by anti-virus scanning for the listener, the message is processed and the verdict returned is printed. If the Cisco appliance is configured to “clean” infected messages, that information is noted. If configured to bounce or drop the messages based on the verdict, that information is printed and the <code>trace</code> command processing stops.</p> <p>Note: This step is skipped if anti-virus scanning is unavailable on the system. If anti-virus scanning is available but has not been enabled with a feature key, that information is also printed in this section.</p> <p>See the <a href="#">Anti-Virus</a>.</p> |
| Content Filters Processing  | <p>All content filters that are enabled on the system are evaluated by the test message at this point. For each filter, the rule is evaluated, and if the end result is “true,” each of the actions in that filter are then performed in sequence. A filter may contain other filters as an action, and the nesting of filters is unlimited. The results of the content filters, processed in order, are printed in this section.</p> <p>See <a href="#">Content Filters</a>.</p>                                                                                                                                                                                                                                                                             |
| Outbreak Filters Processing | <p>This section notes that messages that contain attachments are to bypass the Outbreak Filters feature. If messages are to be processed by Outbreak Filters for the recipient, the message is processed and the evaluation. If the appliance is configured to quarantine, bounce, or drop the messages based on the verdict, that information is printed and the processing stops.</p> <p>See <a href="#">Outbreak Filters</a>.</p>                                                                                                                                                                                                                                                                                                                          |

**Table 40-2**      *Viewing Output When Performing a Trace (continued)*

| trace Command Section | Output                                                                                                                                                                                                                    |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Footer Stamping       | This section notes whether a footer text resource was appended to the message. The name of the text resource is displayed. See <a href="#">Message Disclaimer Stamping, page 21-2</a> in <a href="#">Text Resources</a> . |

**Table 40-2**      **Viewing Output When Performing a Trace (continued)**

| trace Command Section                                                                                                                                                                                    | Output                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Delivery Operations</b>                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                            |
| The following sections note operations that occur when a message is delivered. The <code>trace</code> command prints “Message Enqueued for Delivery” before this section.                                |                                                                                                                                                                                                                                                                                                                            |
| Global Unsubscribe per Domain and per User                                                                                                                                                               | If any recipients you specified as input for the <code>trace</code> command match recipients, recipient domains, or IP addresses listed in the in the Global Unsubscribe feature, any unsubscribed recipient addresses are printed in this section.<br><br>See <a href="#">Configuring Routing and Delivery Features</a> . |
| <b>Final Result</b>                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                            |
| When all processing has been printed, you are prompted with the final result. In the CLI, Answer <b>y</b> to the question, “Would you like to see the resulting message?” to view the resulting message. |                                                                                                                                                                                                                                                                                                                            |

## GUI Example of the Trace Page

**Figure 40-1** *Input for the Trace Page*  
**Trace**

| Message Definition                               |                                                                                                                                       |
|--------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <b>Sender Information</b>                        |                                                                                                                                       |
| Source IP:                                       | <input type="text" value="1.2.3.4"/>                                                                                                  |
| Fully Qualified Domain Name of the Source IP: ?  | <input type="text" value="remotehost.example.com"/>                                                                                   |
| Listener to Trace Behavior on:                   | <input type="text" value="Public (172.22.85.1:25)"/>                                                                                  |
| SenderBase Network Owner ID:                     | <input checked="" type="radio"/> Lookup network owner ID associated with source IP<br><input type="radio"/> Use: <input type="text"/> |
| SenderBase Reputation Score (SBRS):              | <input checked="" type="radio"/> Lookup SBRS associated with source IP<br><input type="radio"/> Use: <input type="text"/>             |
| <b>Envelope Information</b>                      |                                                                                                                                       |
| Envelope Sender:                                 | <input type="text" value="pretend.sender@example.domain"/>                                                                            |
| Envelope Recipients (separated by commas):       | <input type="text" value="admin@ironport.com"/>                                                                                       |
| <b>Message Body</b>                              |                                                                                                                                       |
| Upload Message Body:                             | <input type="text"/> <input type="button" value="Browse..."/>                                                                         |
| Paste Message Body:<br>(If no file is uploaded.) | <input type="text" value="Subject: hello"/><br><input type="text" value="This is a test message."/>                                   |

**Figure 40-2**      **Output for the Trace Page (1 of 2)**  
**Trace**

| Trace Results                                   |                                                 |         |         |
|-------------------------------------------------|-------------------------------------------------|---------|---------|
| Host Access Table Processing (Listener: Public) |                                                 |         |         |
| Matched On:                                     | ALL Sender Group                                |         |         |
| Named Policy:                                   | ACCEPTED                                        |         |         |
| Connection Behavior:                            | ACCEPT                                          |         |         |
| Fully Qualified Domain Name:                    |                                                 |         |         |
| SenderBase Network Owner ID:                    | N/A                                             |         |         |
| SenderBase Reputation Score:                    | N/A                                             |         |         |
| Policy Parameters:                              | Max. Messages Per Connection:                   | 1,000   | Default |
|                                                 | Max. Recipients Per Message:                    | 1,000   | Default |
|                                                 | Max. Message Size:                              | 100M    | Default |
|                                                 | Max. Concurrent Connection From a Single IP:    | 1,000   | Default |
|                                                 | Use TLS:                                        | No      | Default |
|                                                 | Max. Recipients Per Hour:                       | 1000    |         |
|                                                 | Use SenderBase:                                 | Yes     |         |
|                                                 | Use Spam Detection:                             | Yes     |         |
| Use Virus Detection:                            | Yes                                             | Default |         |
| Envelope Sender Processing                      |                                                 |         |         |
| Envelope Sender: pretend.sender@example.domain  |                                                 |         |         |
| Default Domain Processing:                      | No Change                                       |         |         |
| Envelope Recipient Processing                   |                                                 |         |         |
| Envelope Recipient: admin@ironport.com          |                                                 |         |         |
| Default Domain Processing:                      | No Change                                       |         |         |
| Domain Map Processing:                          | No Change                                       |         |         |
| Recipient Access Table Processing:              | Behavior: ACCEPT Matched On: admin@ironport.com |         |         |
| Alias Expansion:                                | No Change                                       |         |         |
| Message Processing                              |                                                 |         |         |
| Assigned Virtual Gateway:                       | None                                            |         |         |
| Assigned Bounce Profile:                        | None                                            |         |         |




|                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Domain Masquerading                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|                                                                    | No changes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Filter Processing                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| skipper                                                            | Skipped (Inactive)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| always_deliver                                                     | Rule: rcpt-to == "@mail.qa": False<br>Rule: rcpt-to == "ironport.com": True<br>Rule: OR: True<br>Action: deliver()                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Mail Policy Processing: Inbound (matched on policy Public Upgrade) |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Message going to:                                                  | admin@ironport.com                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Anti-Spam Processing                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Evaluation:                                                        | Not Spam                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Anti-Virus Processing                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Evaluation:                                                        | No Viruses Detected<br>Elapsed Time: 0.000 sec                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Actions Taken:                                                     | Delivered                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| VOF Processing                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Evaluation:                                                        | No threat detected                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Footer Stamping                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Appended Text Resource:                                            | footer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| DomainKey Signing                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Result of DomainKeys processing:                                   | DomainKeys signing not enabled in this listener's HAT                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Message Delivery (matched on policy Public Upgrade)                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Final Envelope Sender:                                             | pretend.sender@example.domain                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Final Recipients:                                                  | admin@ironport.com                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Final Message:                                                     | <p>Received: from remotehost.example.com (HELO TEST) ([1.2.3.4])<br/>by mail3.example.com with TEST; 21 Jul 2005 14:40:05 -0700<br/>Message-Id: &lt;48q06k\$@Public&gt;<br/>X-Brightmail-Tracker: AAAAAA==<br/>X-BrightmailFiltered: true<br/>X-IronPort-Anti-Spam-Filtered: true<br/>X-IronPort-AV: i="3.95,134,1120460400";<br/>d="scan"; a="0:sNHT0"</p> <p>Subject: hello<br/>Content-Transfer-Encoding: base64<br/>Content-Type: text/plain; charset="utf-8"</p> <p>VGhpcyBpcyBhIHRlc3QgbWVzc2FnZS4KPT09PT09PT09CuODleODg+OCv+ODv+OObP+OBmeOA<br/>guOCj+OBhOOCj+OBhOOAggpUaGlzIGlzIGEgSmFwYW5lc2UgZm9vdGVyCj09PT09PT09PT09PQo=</p> |

Done

“Black hole” listeners allow you to test your message generation systems and to also get a rough measure of receiving performance. Two types of black hole listeners are *queueing* and *non-queueing*.

- For example, in [Figure 40-4](#), you could create a black hole listener “C” to mirror the private listener labeled “B.” A non-queueing version tests the performance path of the system from the groupware client to the groupware server to the appliance. A queueing version tests that same path *and* the appliance’s ability to enqueue messages and prepare them for delivery via SMTP.

**Figure 40-4** *Black Hole Listener for an Enterprise Gateway*

In the following example, the `listenerconfig` command is used to create a black hole queueing listener named `BlackHole_1` on the Management interface. This Host Access Table (HAT) for the listener is then edited to accept connections from the following hosts:

- `yoursystem.example.com`
- `10.1.2.29`
- `badmail.tst`
- `.tst`

**Note**

The final entry, `.tst`, configures the listener so that any host in the `.tst` domain can send email to the listener named `BlackHole_1`.

## Example

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

Currently configured listeners:

1. InboundMail (on PublicNet, 192.168.2.1) SMTP Port 25 Public
2. OutboundMail (on PrivateNet, 192.168.1.1) SMTP Port 25 Private

Choose the operation you want to perform:

- NEW - Create a new listener.
- EDIT - Modify a listener.
- DELETE - Remove a listener.

- SETUP - Change global settings.

[> **new**

Please select the type of listener you want to create.

1. Private
2. Public
3. Blackhole

[2]> **3**

Do you want messages to be queued onto disk? [N]> **y**

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

[> **BlackHole\_1**

Please choose an IP interface for this Listener.

1. Management (192.168.42.42/24: mail3.example.com)
2. PrivateNet (192.168.1.1/24: mail3.example.com)
3. PublicNet (192.168.2.1/24: mail3.example.com)

[1]> **1**

Choose a protocol.

1. SMTP
2. QMQP

[1]> **1**

Please enter the IP port for this listener.

[25]> **25**

Please specify the systems allowed to relay email through the IronPort C60.

Hostnames such as "example.com" are allowed.

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

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

Separate multiple entries with commas.

```
[> yoursystem.example.com, 10.1.2.29, badmail.tst, .tst
```

Do you want to enable rate limiting per host? (Rate limiting defines

the maximum number of recipients per hour you are willing to receive from a remote domain.) [N]> **n**

Default Policy Parameters

=====

Maximum Message Size: 100M

Maximum Number Of Connections From A Single IP: 600

Maximum Number Of Messages Per Connection: 10,000

Maximum Number Of Recipients Per Message: 100,000

Maximum Number Of Recipients Per Hour: Disabled

Use SenderBase for Flow Control: No

Spam Detection Enabled: No

Virus Detection Enabled: Yes

Allow TLS Connections: No

Allow SMTP Authentication: No

Require TLS To Offer SMTP authentication: No

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

Listener BlackHole\_1 created.

Defaults have been set for a Black Hole Queuing listener.

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

Currently configured listeners:

1. BlackHole\_1 (on Management, 192.168.42.42) SMTP Port 25 Black Hole Queuing
2. InboundMail (on PublicNet, 192.168.1.1) SMTP Port 25 Public
3. OutboundMail (on PrivateNet, 192.168.1.1) SMTP Port 25 Private

Choose the operation you want to perform:

- NEW - Create a new listener.
- EDIT - Modify a listener.
- DELETE - Remove a listener.
- SETUP - Change global settings.

[ ]>



#### Note

Remember to issue the `commit` command for these changes to take effect.

After you have configured a black hole queuing listener and modified the HAT to accept connections from your injection system, use your injection system to begin sending email to the appliance. Use the `status`, `status detail`, and `rate` commands to monitor system performance. You can also monitor the system via the Graphical User Interface (GUI). For more information, see:

- [Monitoring Using the CLI, page 34-6](#)
- [Other Tasks in the GUI, page 36-7](#)

## Troubleshooting the Network

If you suspect that the appliance has network connectivity issues, first confirm that the appliance is working properly.

## Testing the Network Connectivity of the Appliance

### Procedure

- Step 1** Connect to the system and log in as the administrator. After successfully logging in, the following messages are displayed:

Last login: *day month date hh:mm:ss* from *IP address*

Copyright (c) 2001-2003, IronPort Systems, Inc.

```
AsyncOS x.x for Cisco
```

```
Welcome to the Cisco Messaging Gateway Appliance(tm)
```

**Step 2** Use the `status` or `status detail` commands.

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

or

```
mail3.example.com> status detail
```

The `status` command returns a subset of the monitored information about email operations. The statistics returned are grouped into two categories: counters and gauges. For complete monitoring information about email operations including rates, use the `status detail` command. Counters provide a running total of various events in the system. For each counter, you can view the total number of events that have occurred since the counter was reset, since the last system reboot, and over the system's lifetime. (For more information, see [Monitoring Using the CLI, page 34-6](#).)

**Step 3** Use the `mailconfig` command to send mail to a known working address.

The `mailconfig` command generates a human-readable file including all configuration settings available to the appliance. Attempt to send the file from the appliance to a known working email address to confirm that the appliance is able to send email over the network.

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

```
Please enter the email address to which you want to send the
configuration file.
```

```
Separate multiple addresses with commas.
```

```
[> user@example.com
```

```
Do you want to include passwords? Please be aware that a configuration without passwords
will fail when reloaded with loadconfig. [N]> y
```

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

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

## Troubleshooting

After you have confirmed that the appliance is active on the network, use the following commands to pinpoint any network problems.

- You can use the `netstat` command to display network connections (both incoming and outgoing), routing tables, and a number of network interface statistics, including the following information:
  - List of active sockets
  - State of network interfaces
  - Contents of routing tables
  - Size of the listen queues
  - Packet traffic information
- You can use the `diagnostic -> network -> flush` command to flush all network related caches.
- You can use the `diagnostic -> network -> arpshow` command to show the system ARP cache.
- You can use the `packetcapture` command to intercept and display TCP/IP and other packets being transmitted or received over a network to which the computer is attached.

To use `packetcapture`, set the network interface and the filter. The filter uses the same format the UNIX `tcpdump` command. Use `start` to begin the packet capture and `stop` to end it. After stopping the capture, you need to use SCP or FTP to download the files from the `/pub/captures` directory. For more information, see [Running a Packet Capture, page 40-31](#).

- Use the `ping` command to a known working host to confirm that the appliance has an active connection on the network and is able to reach specific segments of your network.

The `ping` command allows you to test connectivity to a network host from the appliance.

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

```
Which interface do you want to send the pings from?
```

1. Auto
2. Management (192.168.42.42/24: mail3.example.com)
3. PrivateNet (192.168.1.1/24: mail3.example.com)
4. PublicNet (192.168.2.1/24: mail3.example.com)

```
[1]> 1
```

```
Please enter the host you wish to ping.
```

```
[> anotherhost.example.com
```

```
Press Ctrl-C to stop.
```

```
PING anotherhost.example.com (x.x.x.x): 56 data bytes
```

```

64 bytes from 10.19.0.31: icmp_seq=9 ttl=64 time=0.133 ms
64 bytes from 10.19.0.31: icmp_seq=10 ttl=64 time=0.115 ms
^C
--- anotherhost.example.com ping statistics ---
11 packets transmitted, 11 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.115/0.242/1.421/0.373 ms

```

**Note**

You must use Control-C to end the `ping` command.

- Use the `traceroute` command to test connectivity to a network host from the appliance and debug routing issues with network hops.

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

Which interface do you want to trace from?

1. Auto
2. Management (192.168.42.42/24: mail3.example.com)
3. PrivateNet (192.168.1.1/24: mail3.example.com)
4. PublicNet (192.168.2.1/24: mail3.example.com)

```
[1]> 1
```

Please enter the host to which you want to trace the route.

```
[> 10.1.1.1
```

Press Ctrl-C to stop.

```
traceroute to 10.1.1.1 (10.1.1.1), 64 hops max, 44 byte packets
```

- ```

1  gateway (192.168.0.1)  0.202 ms  0.173 ms  0.161 ms
2  hostname (10.1.1.1)  0.298 ms  0.302 ms  0.291 ms

```

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

- Use the `diagnostic -> network -> smtping` command to test a remote SMTP server.

- Use the `nslookup` command to check the DNS functionality.

The `nslookup` command can confirm that the appliance is able to reach and resolve hostnames and IP addresses from a working DNS (domain name service) server.

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

```
Please enter the host or IP to resolve.
```

```
[> example.com
```

```
Choose the query type:
```

1. A
 2. CNAME
 3. MX
 4. NS
 5. PTR
 6. SOA
 7. TXT
- ```
[1]>
```

```
A=192.0.34.166 TTL=2d
```

**Table 40-3**      *Checking DNS Functionality: Query Types*

| Query Type | Description                                                                                  |
|------------|----------------------------------------------------------------------------------------------|
| A          | the host's Internet address                                                                  |
| CNAME      | the canonical name for an alias                                                              |
| MX         | the mail exchanger                                                                           |
| NS         | the name server for the named zone                                                           |
| PTR        | the hostname if the query is an Internet address, otherwise the pointer to other information |
| SOA        | the domain's "start-of-authority" information                                                |
| TXT        | the text information                                                                         |

- Use the `tophosts` command via the CLI or the GUI, and sort by Active Recipients.

The `tophosts` command returns a list of the top 20 recipient hosts in queue. This command can help you determine if network connectivity problems are isolated to a single host or group of hosts to which you are attempting to send email. (For more information, see “Determining the Make-up of the Mail Queue” on page 49.)

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

```
Sort results by:
```

1. Active Recipients
2. Connections Out
3. Delivered Recipients
4. Soft Bounced Events
5. Hard Bounced Recipients

```
[1]> 1
```

```
Status as of: Mon Nov 18 22:22:23 2003
```

```
ActiveConn.Deliv.SoftHard
```

```
Recipient HostRecipOutRecip.BouncedBounced
```

```
1 aol.com36510255218
2 hotmail.com29071982813
3 yahoo.com13461231119
4 excite.com9838494
5 msn.com8427633 29
```

```
^C
```

- “Drill-down” to use the `hoststatus` command on the top domains listed from the `tophosts` command results.

The `hoststatus` command returns monitoring information about email operations relating to a specific recipient host. DNS information stored in the AsyncOS cache and the last error returned from the recipient host are also given. Data returned is cumulative since the last `resetcounters` command. (For more information, see [Monitoring the Status of a Mail Host, page 34-11](#).)

Using the `hoststatus` command on the top domains can isolate the performance issues with DNS resolution to the either the appliance or the internet. For example, if the `hoststatus` command for the top active recipient host shows many pending outbound connections, then try to determine if that particular host is down or unreachable, or if the appliance cannot connect to all or the majority of hosts.

- Check firewall permissions.

The appliance may need all of the following ports to be opened in order to function properly: ports 20, 21, 22, 23, 25, 53, 80, 123, 443, and 628. (See [Firewall Information](#).)

- Send email from the appliance on your network to `dnscheck@ironport.com`

Send an email from within your network to `dnscheck@ironport.com` to perform basic DNS checks on your system. An auto-responder email will respond with the results and details of the following four tests:

**DNS PTR Record** - Does the IP address of the Envelope From match the PTR record for the domain?

**DNS A Record** - Does the PTR record for the domain match the IP address of the Envelope From?

**HELO match** - Does the domain listed in the SMTP HELO command match the DNS hostname in the Envelope From?

**Mail server accepting delayed bounce messages** - Does the domain listed in the SMTP HELO command have MX records that resolve IP addresses for that domain?

## Troubleshooting the Listener

If you suspect problems with injecting email, use the following strategies:

- Confirm the IP address that you are injecting from, and then use the `listenerconfig` command to check for allowed hosts.

Is the IP address allowed to connect to the listener you have created? Use the `listenerconfig` command to examine the Host Access Table (HAT) for the listener. Use these commands to print the HAT for a listener:

```
listenerconfig -> edit -> listener_number -> hostaccess -> print
```

The HAT can be configured to refuse connections by IP address, block of IP addresses, hostname, or domains. For more information, see “Specifying Hosts that are Allowed to Connect” on page 107.

You can also use the `limits` subcommand to check the maximum number of connections allowed for a listener:

```
listenerconfig -> edit -> listener_number -> limits
```

- On the machine that you are injecting from, use Telnet or FTP to manually connect to the appliance. For example:

```
injection_machine% telnet appliance_name
```

You can also use the `telnet` command within the appliance itself to connect from the listener to the actual appliance:

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

Please select which interface you want to telnet from.

1. Auto

```
2. Management (192.168.42.42/24: mail3.example.com)
3. PrivateNet (192.168.1.1/24: mail3.example.com)
4. PublicNet (192.168.2.1/24: mail3.example.com)
```

```
[1]> 3
```

Enter the remote hostname or IP.

```
[> 193.168.1.1
```

Enter the remote port.

```
[25]> 25
```

Trying 193.168.1.1...

Connected to 193.168.1.1.

Escape character is '^'.

If you cannot connect from one interface to another, you may have issues with the way in which the appliance's Management and Data1 and Data2 interfaces are connected to your network. Ensure that the telnet service is enabled on the target interface if you are attempting to connect using telnet. See [Appendix A, “FTP, SSH, SCP, and Telnet Access”](#) for more information. You can also telnet to port 25 of the listener and enter SMTP commands manually (if you are familiar with the protocol).

- Examine the IronPort text mail logs and injection debug logs to check for receiving errors.

Injection debug logs record the SMTP conversation between the appliance and a specified host connecting to the system. Injection debug logs are useful for troubleshooting communication problems between the appliance and a client initiating a connection from the Internet. The log records all bytes transmitted between the two systems and classifies them as “Sent to” the connecting host or “Received from” the connecting host.

For more information, see [Using Text Mail Logs, page 38-8](#) and [Using Injection Debug Logs, page 38-23](#).

## Troubleshooting Email Delivery From the Appliance

If you suspect problems with delivering email from the appliance, try the following strategies:

- Determine if the problem is domain-specific.

Use the `tophosts` command to get immediate information about the email queue and determine if a particular recipient domain has delivery problems.

Are there problem domains returned when you sort by “Active Recipients?”

When you sort by Connections Out, does any one domain reach the maximum connections specified for a listener? The default maximum number of connections for a listener is 600. The default maximum system-wide number of connections is 10,000 (set by the `deliveryconfig` command). You can examine the maximum number of connections for a listener using the command:

```
listenerconfig -> edit -> listener_number -> limits
```

Are the connections for a listener further limited by the `destconfig` command (either by system maximum or by Virtual Gateway addresses)? Use this command to examine the `destconfig` connection limits:

```
destconfig -> list
```

- Use the `hoststatus` command.

“Drill-down” using the `hoststatus` command on the top domains listed from the results listed by the `tophosts` command.

Is the host available and accepting connections?

Are there problems with one specific MX record mail server for the given host?

The `hoststatus` command reports the last “5XX” status code and description returned by the host if there is a 5XX error (Permanent Negative Completion reply) for the specified host. If the last outgoing TLS connection to the host failed, the `hoststatus` command displays the reason why it failed.

- Configure and/or examine the domain debug, bounce, and text mail logs to check if the recipient host is available.

**Domain debug logs** record the client and server communication during an SMTP conversation between the appliance and a specified recipient host. This log file type can be used to debug issues with specific recipient hosts.

For more information, see [Using Domain Debug Logs, page 38-22](#).

**Bounce logs** record all information pertaining to each bounced recipient.

For more information, see [Using Bounce Logs, page 38-17](#).

**Text mail logs** contain details of email receiving, email delivery and bounces. Status information is also written to the mail log every minute. These logs are a useful source of information to understand delivery of specific messages and to analyze system performance.

For more information, see [Using Text Mail Logs, page 38-8](#).

- Use the `telnet` command to connect from the appliance to the problem domain:

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

```
Please select which interface you want to telnet from.
```

1. Auto
2. Management (192.168.42.42/24: mail3.example.com)
3. PrivateNet (192.168.1.1/24: mail3.example.com)
4. PublicNet (192.168.2.1/24: mail3.example.com)

```
[1]> 1
```

Enter the remote hostname or IP.

```
[> problemdomain.net
```

Enter the remote port.

```
[25]> 25
```

- You can use the `tlsverify` command to establish an outbound TLS connection on demand and debug any TLS connection issues concerning a destination domain. To create the connection, specify the domain to verify against and the destination host. AsyncOS checks the TLS connection based on the Required (Verify) TLS setting.

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

Enter the TLS domain to verify against:

```
[> example.com
```

Enter the destination host to connect to. Append the port (example.com:26) if you are not connecting on port 25:

```
[example.com]> mxe.example.com:25
```

Connecting to 1.1.1.1 on port 25.

Connected to 1.1.1.1 from interface 10.10.10.10.

Checking TLS connection.

TLS connection established: protocol TLSv1, cipher RC4-SHA.

Verifying peer certificate.

Verifying certificate common name mxe.example.com.

TLS certificate match mxe.example.com

TLS certificate verified.

TLS connection to 1.1.1.1 succeeded.

```
TLS successfully connected to mx.example.com.
```

```
TLS verification completed.
```

## Troubleshooting Performance

If you suspect that there are performance problems with the appliance, utilize the following strategies:

- Use the `rate` and `hostrate` commands to check the current system activity.

The `rate` command returns real-time monitoring information about email operations. For more information, see [Displaying Real-time Activity, page 34-16](#).

The `hostrate` command returns real-time monitoring information for a specific host.

- Use the `status` command to cross-check the historical rates to check for degradation.
- Use the `status detail` command to check the RAM utilization.

You can use the `status detail` command to quickly see the system's RAM, CPU, and Disk I/O utilization.



### Note

RAM utilization should always be less than 45%. If RAM utilization exceeds 45%, then, the appliance will enter "resource conservation mode;" it initiates a "back-off" algorithm to prevent over-subscription of resources and sends out the following email alert:

```
This system (hostname: hostname) has entered a 'resource conservation' mode in order to prevent the rapid depletion of critical system resources.
```

```
RAM utilization for this system has exceeded the resource conservation threshold of 45%. The allowed injection rate for this system will be gradually decreased as RAM utilization approaches 60%.
```

This situation occurs only with an aggressive injection with poor deliverability facilities. If you encounter RAM utilization exceeding 45%, check the number of messages in the queue and see if a particular domain is down or unavailable for delivery (via the `hoststatus` or `hostrate` commands). Also check the status of the system and ensure that delivery is not suspended. If after stopping the injection you continue to experience a high RAM utilization, contact Cisco Customer Support. See [Cisco Customer Support, page 1-9](#).

- Is the problem specific to one domain?

Use the `tophosts` command to get immediate information about the email queue and determine if a particular recipient domain has delivery problems.

Check the size of the queue. You can delete, bounce, suspend, or redirect messages in the email queue to manage its size, or to deal with recipients to a specific, problematic domain. For more information, see [Managing the Email Queue, page 34-22](#). Use these commands:

- `deleterecipients`
- `bouncerecipients`

- `redirectrecipients`
- `suspenddel / resumedel`
- `suspendlistener / resumelister`

Use the `tophosts` command to check the number of soft and hard bounces. Sort by “Soft Bounced Events” (option 4) or “Hard Bounced Recipients” (option 5). If the performance for a particular domain is problematic, use the commands above to manage the delivery to that domain.

## Responding to Alerts

- [Alert: Battery Relearn Timed Out \(RAID Event\) on C380 or C680 Hardware, page 40-27](#)
- [Troubleshooting Alerts That Miscellaneous Disk Usage is Approaching the Quota, page 40-27](#)

### Alert: Battery Relearn Timed Out (RAID Event) on C380 or C680 Hardware

**Problem** You have C380 or C680 hardware and you receive a “Battery Relearn Timed Out” (RAID event) alert.

**Solution** This alert may or may not indicate a problem. The battery relearn timeout, in itself, does not mean there is any problem with the RAID controller. The controller can recover in the subsequent relearn. Please monitor your email for any other RAID alerts for the next 48 hours, to ensure that this is not the side-effect of any other problem. If you do not see any other RAID-related alerts from the system, then you can safely ignore this alert.

### Troubleshooting Alerts That Miscellaneous Disk Usage is Approaching the Quota

**Problem** You receive an alert that Miscellaneous disk usage is approaching its quota.

**Solution** You can either increase the quota or delete files. See [Managing Disk Space for the Miscellaneous Quota, page 33-17](#).

## Remotely Resetting Appliance Power

If the appliance requires a hard reset, you can reboot the appliance chassis remotely using a third-party Intelligent Platform Management Interface (IPMI) tool.

#### Restrictions

- Remote power management is available only on certain hardware.  
For specifics, see [Enabling Remote Power Management, page 33-29](#).
- If you want be able to use this feature, you must enable it in advance, before you need to use it.  
For details, see [Enabling Remote Power Management, page 33-29](#).
- Only the following IPMI commands are supported:

`status, on, off, cycle, reset, diag, soft`



Issuing unsupported commands will produce an “insufficient privileges” error.

**Before You Begin**

- Obtain and set up a utility that can manage devices using IPMI version 2.0.
- Understand how to use the supported IPMI commands. See the documentation for your IPMI tool.

**Procedure**

- 
- Step 1** Use IPMI to issue a supported power-cycling command to the IP address assigned to the Remote Power Management port, which you configured earlier, along with the required credentials.

For example, from a UNIX-type machine with IPMI support, you might issue the command:

```
ipmitool -I lan -H 192.0.2.1 -U remoteresetuser -P password chassis power reset
```

where `192.0.2.1` is the IP address assigned to the Remote Power Management port and `remoteresetuser` and `password` are the credentials that you entered while enabling this feature.

- Step 2** Wait at least five minutes for the appliance to reboot.
- 

## Working with Technical Support

- [Technical Support for Virtual Appliances, page 40-28](#)
- [Opening or Updating a Support Case From the Appliance, page 40-28](#)
- [Enabling Remote Access for Cisco Technical Support Personnel, page 40-29](#)
- [Running a Packet Capture, page 40-31](#)

## Technical Support for Virtual Appliances

Requirements for getting technical support for your virtual appliance are described in the *Cisco Content Security Virtual Appliance Installation Guide* available from <http://www.cisco.com/c/en/us/support/security/email-security-appliance/products-installation-guides-list.html>.

## Opening or Updating a Support Case From the Appliance

**Before You Begin**

- If your issue is urgent, do not use this method. Instead, contact support using one of the other methods listed in [Cisco Customer Support, page 1-9](#).

Use the following procedure only for issues such as a request for information or a problem for which you have a workaround, but would like an alternate solution.

- Consider other options for getting help:
  - [Knowledge Base, page 1-9](#)
  - [Cisco Support Community, page 1-9](#)

- To access Cisco technical support directly from the appliance, your Cisco.com user ID must be associated with your service agreement contract for this appliance. To view a list of service contracts that are currently associated with your Cisco.com profile, visit the Cisco.com Profile Manager at [https://tools.cisco.com/RPFA/profile/profile\\_management.do](https://tools.cisco.com/RPFA/profile/profile_management.do). If you do not have a Cisco.com user ID, register to get one. See [Registering for a Cisco Account, page 1-10](#).

Be sure to save your Cisco.com user ID and support contract ID in a safe location.

- When you open a support case using this procedure, the appliance configuration file is sent to Cisco Customer Support. If you do not want to send the appliance configuration, you can contact Customer Support using a different method.
- In cluster configurations, support requests and their saved values are machine-specific.
- The appliance must be connected to the internet and able to send email.
- If you are sending information about an existing case, make sure you have the case number.

#### Procedure

- 
- Step 1** Sign in to the appliance.
  - Step 2** Choose **Help and Support > Contact Technical Support**.
  - Step 3** Complete the form.
  - Step 4** Click **Send**.



---

**Note** CCO User IDs and the last-entered Contract ID are saved on the appliance for future use.

---

## Enabling Remote Access for Cisco Technical Support Personnel

Only Cisco Customer Assistance can access your appliance using these methods.

- [Enabling Remote Access to Appliances With an Internet Connection, page 40-29](#)
- [Enabling Remote Access to Appliances Without a Direct Internet Connection, page 40-30](#)
- [Disabling a Tech Support Tunnel, page 40-31](#)
- [Disabling Remote Access, page 40-31](#)
- [Checking the Status of the Support Connection, page 40-31](#)

## Enabling Remote Access to Appliances With an Internet Connection

Support accesses the appliance through an SSH tunnel that this procedure creates between the appliance and the `upgrades.ironport.com` server.

#### Before You Begin

Identify a port that can be reached from the internet. The default is port 25, which will work in most environments because the system also requires general access over that port in order to send email messages. Connections over this port are allowed in most firewall configurations.

**Procedure**

- 
- Step 1** Log in to the appliance.
- Step 2** From the top right side of the GUI window, choose **Help and Support > Remote Access**.
- Step 3** Click **Enable**.
- Step 4** Enter information:

| Option                    | Description                                                                                                                                                                                |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Customer Support Password | This interim password and the appliance serial number (for physical appliances) or VLN (for virtual appliances) will be used to generate a password for Support access.                    |
| Secure Tunnel             | Select the check box to use a secure tunnel for the remote access connection.<br><br>Enter a port for the connection.<br><br>The default is port 25, which will work in most environments. |

- Step 5** Click **Submit**.
- 

**What To Do Next**

When remote access for support personnel is no longer required, see [Disabling a Tech Support Tunnel, page 40-31](#).

## Enabling Remote Access to Appliances Without a Direct Internet Connection

For appliances without a direct internet connection, access is made through a second appliance that is connected to the internet.

**Before You Begin**

- The appliance must be able to connect on port 22 to a second appliance that is connected to the internet.
- On the appliance with the internet connection, follow the procedure in [Enabling Remote Access to Appliances With an Internet Connection, page 40-29](#) to create a support tunnel to that appliance.

**Procedure**

- 
- Step 1** From the command-line interface of the appliance requiring support, enter the `techsupport` command.
- Step 2** Enter `sshaccess`.
- Step 3** Follow the prompts.
- 

**What To Do Next**

When remote access for support personnel is no longer required, see the following:

- [Disabling Remote Access, page 40-31](#)
- [Disabling a Tech Support Tunnel, page 40-31](#)

## Disabling a Tech Support Tunnel

An enabled techsupport tunnel remains connected to `upgrades.ironport.com` for 7 days. After that time, established connections will not be disconnected but will be unable to re-attach to the tunnel once disconnected.

To disable the tunnel manually:

### Procedure

- 
- |               |                                                                                                |
|---------------|------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Log in to the appliance.                                                                       |
| <b>Step 2</b> | From the top right side of the GUI window, choose <b>Help and Support &gt; Remote Access</b> . |
| <b>Step 3</b> | Click <b>Disable</b> .                                                                         |
- 

## Disabling Remote Access

A remote access account that you create using the `techsupport` command remains active until you deactivate it.

### Procedure

- 
- |               |                                                                              |
|---------------|------------------------------------------------------------------------------|
| <b>Step 1</b> | From the command-line interface, enter the <code>techsupport</code> command. |
| <b>Step 2</b> | Enter <code>sshaccess</code> .                                               |
| <b>Step 3</b> | Enter <code>disable</code> .                                                 |
- 

## Checking the Status of the Support Connection

### Procedure

- 
- |               |                                                                              |
|---------------|------------------------------------------------------------------------------|
| <b>Step 1</b> | From the command-line interface, enter the <code>techsupport</code> command. |
| <b>Step 2</b> | Enter <code>status</code> .                                                  |
- 

## Running a Packet Capture

Packet Capture allows support personnel to see the TCP/IP data and other packets going into and out of the appliance. This allows Support to debug the network setup and to discover what network traffic is reaching the appliance or leaving the appliance.

## Procedure

**Step 1** Choose **Help and Support > Packet Capture**.

**Step 2** Specify packet capture settings:

a. In the **Packet Capture Settings** section, click **Edit Settings**.

b. (Optional) Enter duration, limits, and filters for the packet capture.

Your Support representative may give you guidance on these settings.

If you enter a capture duration without specifying a unit of time, AsyncOS uses seconds by default.

In the Filters section:

- Custom filters can use any syntax supported by the UNIX `tcpdump` command, such as `host 10.10.10.10 && port 80`.
- The client IP is the IP address of the machine connecting to the appliance, such as a mail client sending messages through the Email Security appliance.
- The server IP is the IP address of the machine to which the appliance is connecting, such as an Exchange server to which the appliance is delivering messages.

You can use the client and server IP addresses to track traffic between a specific client and a specific server, with the Email Security appliance in the middle.

c. Click **Submit**.

**Step 3** Click **Start Capture**.

- Only one capture may be running at a time.
- When a packet capture is running, the Packet Capture page shows the status of the capture in progress by showing the current statistics, such as file size and time elapsed.
- The GUI only displays packet captures started in the GUI, not from the CLI. Similarly, the CLI only displays the status of a current packet capture run started in the CLI.
- The packet capture file is split into ten parts. If the file reaches the maximum size limit before the packet capture ends, the oldest part of the file is deleted (the data is discarded) and a new part starts with the current packet capture data. Only 1/10 of the packet capture file is discarded at a time.
- A running capture started in the GUI is preserved between sessions. (A running capture started in the CLI stops when the session ends.)

**Step 4** Allow the capture to run for the specified duration, or, if you have let the capture run indefinitely, manually stop the capture by clicking **Stop Capture**.

**Step 5** Access the packet capture file:

- Click the file in the **Manage Packet Capture Files** list and click **Download File**.
- Use FTP or SCP to access the file in the `captures` subdirectory on the appliance.

## What To Do Next

Make the file available to Support:

- If you allow remote access to your appliance, technicians can access the packet capture files using FTP or SCP. See [Enabling Remote Access for Cisco Technical Support Personnel](#), page 40-29.
- Email the file to Support.



# Optimizing the Appliance for Outbound Mail Delivery Using D-Mode

- [Feature Summary: D-Mode for Optimized Outbound Delivery](#), page 41-1
- [Setting Up the Appliance for Optimized Outbound Mail Delivery](#), page 41-3
- [Sending Bulk Mail Using IronPort Mail Merge \(IPMM\)](#), page 41-4

## Feature Summary: D-Mode for Optimized Outbound Delivery

D-Mode is a feature key-enabled feature that optimizes certain Email Security appliances for outbound email delivery. Features specific to inbound mail handling are disabled in D-Mode.

- [Features Unique to D-Mode-Enabled Appliances](#), page 41-1
- [Standard Features Disabled in D-Mode-Enabled Appliances](#), page 41-2
- [Standard Features Applicable to D-Mode-Enabled Appliances](#), page 41-2

## Features Unique to D-Mode-Enabled Appliances

- 256 Virtual Gateway Addresses - The Cisco Virtual Gateway technology allows you to configure enterprise mail gateways for all domains you host — with distinct IP addresses, hostname and domains — and create separate corporate email policy enforcement and anti-spam strategies for those domains, while hosted within the same physical appliance. See information about “Customizing Listeners” in [Chapter 5, “Configuring the Gateway to Receive Email.”](#)
- IronPort Mail Merge (IPMM) - IronPort Mail Merge (IPMM) removes the burden of generating individual personalized messages from customer systems. By removing the need to generate thousands of individual messages and transmit them between message generating systems and the email gateway, users benefit from the decreased load on their systems and increased throughput of email delivery. For more information, see [Sending Bulk Mail Using IronPort Mail Merge \(IPMM\)](#), page 41-4.
- Resource-conserving bounce setting - You can configure D-Mode-enabled appliances to detect potential blocked destinations and bounce all messages bound for that destination. For more information, see [Configuring Resource-Conserving Bounce Settings](#), page 41-3.
- Enhanced performance for outbound delivery

## Standard Features Disabled in D-Mode-Enabled Appliances

- IronPort anti-spam scanning and on or off box spam quarantining — Because anti-spam scanning pertains mostly to incoming mail, the IronPort Anti-Spam scanning engine is disabled. The Anti-Spam chapter is, therefore, not applicable.
- Outbreak Filters — Because the Outbreak Filters feature is used to quarantine incoming mail, this feature is disabled on D-Mode-enabled appliances. Information in the Outbreak Filters chapter is, therefore, not applicable.
- SenderBase Network Participation capabilities — Because SenderBase Network Participation reports information about incoming mail, this feature is disabled on D-Mode-enabled appliances. Information about SenderBase Network Participation is, therefore, not applicable.
- Reporting — Reporting is limited. Some reports are not available, and the reporting that does occur is set to run at a very limited level for performance reasons.



### Note

The totals shown in the Email Security Monitor Overview report for D-Mode-enabled appliances may erroneously include spam and suspect spam counts, even though these features are disabled on D-Mode-enabled appliances.

- RSA Data Loss Prevention — RSA DLP scanning for outgoing messages is disabled on D-Mode-enabled appliances.

## Standard Features Applicable to D-Mode-Enabled Appliances

**Table 41-1 AsyncOS Features Included in D-Mode Enabled Appliances**

| Feature                  | More Information                                                                                                                                                                                                                                                                                                                                        |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Anti-virus scanning      | See <a href="#">Chapter 12, “Anti-Virus.”</a>                                                                                                                                                                                                                                                                                                           |
| Domain Key signing       | DKIM/Domain Keys is a method for verifying authenticity of email based on a signing key used by the sender. See <a href="#">Chapter 20, “Email Authentication.”</a>                                                                                                                                                                                     |
| Centralized management   | See <a href="#">Chapter 39, “Centralized Management Using Clusters.”</a>                                                                                                                                                                                                                                                                                |
| Delivery throttling      | For each domain, you can assign a maximum number of connections and recipients that will never be exceeded by the system in a given time period. This “good neighbor” table is defined through the <code>destconfig</code> command.<br><br>For more information, see <a href="#">Controlling Email Delivery Using Destination Controls, page 24-42.</a> |
| Bounce Verification      | Verify the authenticity of bounce messages. See <a href="#">Bounce Verification, page 24-51.</a>                                                                                                                                                                                                                                                        |
| Delegated administration | See <a href="#">Chapter 32, “Distributing Administrative Tasks.”</a>                                                                                                                                                                                                                                                                                    |
| Trace (debug)            | See <a href="#">Debugging Mail Flow Using Test Messages: Trace, page 40-1.</a>                                                                                                                                                                                                                                                                          |

**Table 41-1 AsyncOS Features Included in D-Mode Enabled Appliances (continued)**

| Feature                    | More Information                                                                                                                                          |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| VLAN, NIC-pairing          | See <a href="#">Chapter 37, “Advanced Network Configuration.”</a>                                                                                         |
| Optional Anti-virus engine | You can add optional anti-virus scanning to ensure the integrity of your outbound messages. See <a href="#">Anti-Virus Scanning Overview, page 12-1</a> . |

## Setting Up the Appliance for Optimized Outbound Mail Delivery

### Procedure

- Step 1** Apply the provided feature key. You will need to apply the key to your Cisco Email Security appliance *prior to running the system setup wizard* (prior to configuring the appliance). Apply the key via the System Administration > Feature Key page or by issuing the `featurekey` command in the CLI.



**Note** The preceding feature keys include a sample 30 day Sophos or McAfee Anti-Virus license you can use to test anti-virus scanning on outbound mail.

- Step 2** Reboot the appliance.

- Step 3** Run the system setup wizard (GUI or CLI) and configure your appliance.

Please keep in mind that appliances that are optimized for outbound delivery do not include anti-spam scanning or the Outbreak Filters feature. (Please ignore these chapters.)



**Note** In a clustered environment, you cannot combine appliances that are configured with the D-Mode feature key with AsyncOS appliances that are not configured with the delivery performance package.

## Configuring Resource-Conserving Bounce Settings

Once the appliance is configured for optimized outbound mail delivery, you can configure the system to detect potential delivery problems and bounce all messages for a destination.



**Note** Using this setting will bounce all messages in the queue for a destination domain that is deemed undeliverable. You will need to re-send the message once the delivery issues have been resolved.



## Example of Enabling Resource-Conserving Bounce Settings

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

```
Choose the operation you want to perform:
```

- NEW - Create a new profile.
- EDIT - Modify a profile.
- DELETE - Remove a profile.
- SETUP - Configure global bounce settings.

```
[]> setup
```

```
Do you want to bounce all enqueued messages bound for a domain if the host is down? [N]>
y
```

When using this feature, a host is considered “down” after at least 10 consecutive connection attempts fail. AsyncOS scans for down hosts every 15 minutes, so it is possible that more than 10 attempts will be made before the queue is cleared.

## Sending Bulk Mail Using IronPort Mail Merge (IPMM)



### Note

---

IronPort Mail Merge is available only on appliances that are D-Mode-enabled.

---

## Overview of IronPort Mail Merge

IronPort Mail Merge removes the burden of generating individual personalized messages from customer systems. It removes the need to generate thousands of individual messages and transmit them between message generating systems and the email gateway, resulting in decreased load on your systems and increased throughput of email delivery.

With IPMM, a single message body is created with variables representing locations in the message to be replaced for personalization. For each individual message recipient, only the recipient email address and the variable substitutions need to be transmitted to the email gateway. In addition, IPMM can be used to send certain recipients specific “parts” of the message body, while excluding certain parts from others recipients. (For example, suppose you needed to include a different copyright statements at the end of your messages to recipients in two different countries.)

## Benefits of the Mail Merge Function

- Ease of use for the mail administrator. The complexities of creating personalized messages for each recipient are removed, as IPMM provides variable substitution and an abstracted interface in many common languages.
- Reduced load on message generation systems. By requiring one copy of the message body and a table of required substitutions, most of the message generation “work” is off-loaded from message generation systems and moved to the appliance that is configured for optimized outbound mail delivery.
- Increased delivery throughput. By reducing the resources necessary to accept and queue thousands of incoming messages, the appliance can significantly increase out-bound delivery performance.
- Queue storage efficiency. By storing less information for each message recipient, users can achieve orders-of- magnitude, better use of queue storage on the D-Mode enabled appliance.

## Using Mail Merge

### SMTP Injection

IPMM extends SMTP as the transport protocol. There is no special configuration that needs to be made to the appliance. (By default, IPMM can be enabled for private listeners and disabled for public listeners on the D-Mode-enabled appliance.) However, if you are not currently using SMTP as your injection protocol, you must create a new private listener that utilizes SMTP through the D-Mode enabled appliance interface.

Use the `setipmm` subcommand of `listenerconfig` to enable IPMM on the listener. For more information, see [Chapter 5, “Configuring the Gateway to Receive Email.”](#)

IPMM modifies SMTP by altering two commands — `MAIL FROM` and `DATA` — and adding another: `XDFN`. The `MAIL FROM` command is replaced with `XMRG FROM` and, the `DATA` command is replaced with `XPRT`.

To generate a Mail Merge message, the commands used to generate the message need to be issued in a particular sequence.

1. The initial EHLO statement, identifying the sending host.
2. Each message starts with an `XMRG FROM:` statement, indicating the sender address.
3. Each recipient is then defined:
  - One or more `XDFN` variable allocation statements are made, including defining the parts (`XDFN *PART=1,2,3...`), and any other recipient specific variables.
  - The recipient email address is defined with the `RCPT TO:` statement. Any variable allocations prior to the `RCPT TO:`, but after the `XMRG FROM`, or `RCPT TO` command, will be mapped to this recipient email address.
4. Each part is defined using the `XPRT n` command, with each part terminated by a period (.) character similar to the `DATA` command. The last part is defined by the `XPRT n LAST` command.

## Variable Substitution

Any part of the message body, including message headers, can contain variables for substitution. Variables can appear in HTML messages, as well. Variables are user-defined and must begin with the ampersand (&) character and end with the semi-colon character (;). Variable names beginning with an asterisk (\*) are reserved and cannot be used.

## Reserved Variables

IPMM contains five special “reserved” variables that are predefined.

**Table 41-2** *IPMM: Reserved Variables*

|        |                                                                                                                                                                                                                                                  |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| *FROM  | The reserved variable *FROM is derived from the “Envelope From” parameter. The “Envelope From” parameter is set by the “XMRG FROM:” command.                                                                                                     |
| *TO    | The reserved variable *TO is derived from the envelope recipient value, as set by the “RCPT TO:” command.                                                                                                                                        |
| *PARTS | The reserved variable *PARTS holds a comma separated list of parts. It is set prior to defining a recipient with the “RCPT TO:” and determines which of the “XPRT n” message body blocks a given user will receive.                              |
| *DATE  | The reserved variable *DATE is replaced with the current date stamp.                                                                                                                                                                             |
| *DK    | The reserved variable *DK is used to specify a DomainKeys Signing profile (this profile must already exist in AsyncOS). For more information about creating DomainKeys Signing profiles, see <a href="#">Chapter 20, “Email Authentication.”</a> |

## Example Message #1

The following example message body (including headers) contains four distinct variables and five substitution locations that will be replaced in the final message. Note that the same variable may be used more than once in the message body. Also, the reserved variable &\*TO; is used, which will be replaced with the recipient email address. This reserved variable does not need to be passed in as a separate variable. The variables in the example appear in bold.

```
From: Mr.Spacely <spacely@example.com>

To: &first_name;&last_name;&*TO;

Subject: Thanks for Being an Example.Com Customer
```

Dear &first\_name; ,

Thank you for purchasing a &color; sprocket.

This message needs only be injected once into the appliance. For each recipient, the following additional information is required:

- A recipient email address

- Name-value pairs for the variable substitution

## Part Assembly

Where SMTP uses a single `DATA` command for each message body, IPMM uses one or many `XPR` commands to comprise a message. Parts are assembled based upon the order specified per-recipient. Each recipient can receive any or all of the message parts. Parts can be assembled in any order.

The special variable `*PARTS` holds a comma separated list of parts.

For example, the following example message contains two parts.

The first part contains the message headers and some of the message body. The second part contains an offer that can be variably included for specific customers.

### Example Message #2, Part 1

```
From: Mr. Spacely <spacely@example.com>

To: &first_name; &last_name; &*TO;

Subject: Thanks for Being an Example.Com Customer

Dear &first_name;,

Thank you for purchasing a &color; sprocket.
```

### Example Message #2, Part 2

```
Please accept our offer for 10% off your next sprocket purchase.
```

The message parts need only be injected once into the appliance. In this case, each recipient requires the following additional information:

- The ordered list of parts to be included in the final message
- A recipient email address
- Name value pairs for the variable substitution

## IPMM and DomainKeys Signing

IPMM does support DomainKeys Signing. Use the `*DK` reserved variable to specify a DomainKeys profile. For example:

```
XDFN first_name="Jane" last_name="User" color="red" *PARTS=1,2 *DK=mass_mailing_1
```

In this example, “mail\_mailing\_1” is the name of a previously configured DomainKeys profile.

## Command Descriptions

When a client injects IPMM messages to the listener, it uses extended SMTP with the following key commands.

### XMRG FROM

Syntax:

```
XMRG FROM: <sender email address>
```

This command replaces the SMTP `MAIL FROM:` command and indicates that what follows is an IPMM message. An IPMM job is initiated with the `XMRG FROM:` command.

### XDFN

Syntax:

```
XDFN <KEY=VALUE> [KEY=VALUE]
```

The `XDFN` command sets the per-recipient metadata. Note that key-value pairs can optionally be enclosed in angle brackets or square brackets.

`*PARTS` is a special reserved variable that indicates the index number as defined by the `XPRT` command (described below). The `*PARTS` variable is split as a comma-delimited list of integers. The integers match the body parts to be sent as defined by the `XPRT` commands. The other reserved variables are: `*FROM`, `*TO`, and `*DATE`.

### XPRT

Syntax:

```
XPRT index_number LAST
```

```
Message
```

```
.
```

The `XPRT` command replaces the SMTP `DATA` command. The command accepts the transfer of the message part after the command is issued. The command is completed with a single period on a line followed by a return (which is the same way an SMTP `DATA` command is completed).

The special keyword **LAST** indicates the end of the mail merge job and must be used to specify the final part that will be injected.

After the `LAST` keyword is used, the message is queued, and delivery begins.

## Notes on Defining Variables

- When you define variables with the `XDFN` command, note that the actual command line cannot exceed the physical limit of the system. In the case of the D-Mode-enabled appliance, this limit is 4 kilobytes per line. Other host systems may have lower thresholds. Use caution when defining multiple variables on very large lines.

- You can escape special characters using the forward slash “/” character when defining variables key-value pairs. This is useful if your message body contains HTML character entities that might be mistakenly replaced with variable definitions. (For example, the character entity `&trade;` defines the HTML character entity for a trademark character. If you created the command `XDFN trade=foo` and then created a IPMM message containing the HTML character entity “`&trade;`” the assembled message would contain the variable substitution (“`foo`”) instead of the trademark character. The same concept is true for the ampersand character “`&`” which is sometimes used in URLs containing GET commands.

## Example IPMM Conversation

The following is an example IPMM conversation of Example Message #2 (shown above). The message will be sent to two recipients in this example: “Jane User” and “Joe User.”

In this example, the type in **bold** represents what you would type in a manual SMTP conversation with the D-Mode-enabled appliance, type in *monospaced type* represents the responses from the SMTP server, and *italic type* represents comments or variables.

A connection is established:

220 ESMTP

**EHLO** foo

*250-ehlo responses from the listener enabled for IPMM*

The conversation is started:

**XMRG FROM:** <user@domain.com> *[Note: This replaces the MAIL FROM: SMTP command.]*

250 OK

Variables and parts are set for each recipient:

**XDFN first\_name="Jane" last\_name="User" color="red" \*PARTS=1,2**

*[Note: This line defines three variables (first\_name, last\_name, and color) and then uses the \*PARTS reserved variable to define that the next recipient defined will receive message parts numbers 1 and 2.]*

250 OK

**RCPT TO:** <jane@company.com>

250 recipient <jane@company.com> ok

**XDFN first\_name="Joe" last\_name="User" color="black" \*PARTS=1**

*[Note: This line defines three variables (first\_name, last\_name, and color) and then uses the \*PARTS reserved variable to define that the next recipient defined will receive message parts numbers 1 only.]*

```
RCPT TO:<joe@company1.com>
```

```
250 recipient <joe@company1.com> ok
```

Next, part 1 is transmitted:

```
XPRT 1 [Note: This replaces the DATA SMTP command.]
```

```
354 OK, send part
```

```
From: Mr. Spacely <spacely@example.com>
```

```
To: &first_name; &last_name; &*TO;
```

```
Subject: Thanks for Being an Example.Com Customer
```

```
&*DATE;
```

```
Dear &first_name;,
```

```
Thank you for purchasing a &color; sprocket.
```

```
.
```

And then part 2 is transmitted. Note that the `LAST` keyword is used to identify Part 2 as the final part to assemble:

```
XPRT 2 LAST
```

```
Please accept our offer for 10% off your next sprocket purchase.
```

```
.
```

```
250 Ok, mailmerge message enqueued
```

The “250 Ok, mailmerge message queued” notes that the message has been accepted.

Based on this example, recipient Jane User will receive this message:

```
From: Mr. Spacely <spacely@example.com>
```

```
To: Jane User <jane@company.com>
```

```
Subject: Thanks for Being an Example.Com Customer
```

*message date*

Dear Jane,

Thank you for purchasing a red sprocket.

Please accept our offer for 10% off your next sprocket purchase.

Recipient Joe User will receive this message:

From: Mr. Spacely <spacely@example.com>

To: Joe User <joe@company1.com>

Subject: Thanks for Being an Example.Com Customer

*message date*

Dear Joe,

Thank you for purchasing a black sprocket.

## Example Code

Cisco has created libraries in common programming languages to abstract the task of injecting IPMM messages into the appliance listener enabled for IPMM. Contact Cisco Customer Support for examples of how to use the IPMM library. The code is commented extensively to explain its syntax.







## Centralizing Services on a Cisco Content Security Management Appliance

---

- [Overview of Cisco Content Security Management Appliance Services, page 42-1](#)
- [Network Planning, page 42-2](#)
- [Working with an External Spam Quarantine, page 42-2](#)
- [About Centralizing Policy, Virus, and Outbreak Quarantines, page 42-5](#)
- [Configuring Centralized Reporting, page 42-10](#)
- [Configuring Centralized Message Tracking, page 42-11](#)
- [Using Centralized Services, page 42-11](#)

### Overview of Cisco Content Security Management Appliance Services

The Cisco Content Security Management appliance (M-Series appliance) is an external or “off box” location that provides a single interface to certain services on multiple Email Security appliances.

The Security Management appliance includes the following features:

- External spam quarantine. Holds spam and suspected spam messages for end users, and allow end users and administrators to review messages that are flagged as spam before making a final determination.
- Centralized policy, virus, and outbreak quarantines. Provide a single location behind the firewall to store and manage messages quarantined by anti-virus scanning, outbreak filters, and policies.
- Centralized reporting. Run reports on aggregated data from multiple Email Security appliances.
- Centralized tracking. Track email messages that traverse multiple Email Security appliances.


For complete information about configuring and using your Cisco Content Security Management appliance, see the *Cisco Content Security Management Appliance User Guide*.

## Network Planning

The Cisco Content Security Management appliance lets you separate the end-user interfaces (such as mail applications) from the more secure gateway systems residing in your various DMZs. Using a two-layer firewall can provide you with flexibility in network planning so that end users do not connect directly to the outer DMZ.

Figure 42-1 shows a typical network configuration incorporating the Security Management appliance and multiple DMZs.

**Figure 42-1** Typical Network Configuration with Cisco Content Security Management Appliance



Large corporate data centers can share one Security Management appliance which acts as an external spam quarantine for one or more Email Security appliances. Meanwhile, remote offices can maintain local spam quarantines on Email Security appliances for local use.

## Working with an External Spam Quarantine

- [Mail Flow and the External Spam Quarantine, page 42-2](#)
- [Migrating from a Local Spam Quarantine to an External Quarantine, page 42-3](#)
- [Enabling an External Spam Quarantine and External Safelist/Blocklist, page 42-3](#)
- [Disabling the Local Spam Quarantine to Activate the External Quarantine, page 42-4](#)
- [Troubleshooting an External Spam Quarantine, page 42-5](#)

## Mail Flow and the External Spam Quarantine

If your network is configured as described in Figure 42-1, incoming mail from the Internet is received by appliances in the outer DMZ. Clean mail is sent along to the mail transfer agent (MTA) (groupware) in the inner DMZ and eventually to the end users within the corporate network.

Spam and suspected spam (depending on your mail flow policy settings) is sent to the spam quarantine on the Security Management appliance. End users may then access the quarantine and elect to delete spam and release messages that they would like to have delivered to themselves. Messages remaining in the spam quarantine are automatically deleted after a configurable amount of time.

Messages that are released from the external quarantine on the Security Management appliance are returned to the originating Email Security appliance for delivery. These messages do not normally pass through the following processes before delivery: HAT and other policy or scanning settings, RAT, domain exceptions, aliasing, incoming filters, masquerading, bounce verification, and the work queue.

An Email Security appliance that is configured to send mail to a Security Management appliance will automatically expect to receive mail released from the Security Management appliance and will not reprocess those messages when they are received back. For this to work, the IP address of the Security Management appliance must not change. If the IP address of the Security Management appliance changes, the receiving Email Security appliance will process the message as it would any other incoming message. You should always use the same IP address for receiving and delivery on the Security Management appliance.

The Security Management appliance accepts mail for quarantining from the IP addresses specified in the spam quarantine settings. To configure the spam quarantine on the Security Management appliance, see the *Cisco Content Security Management Appliance User Guide*.

Mail released by the Security Management appliance is delivered to the primary and secondary hosts (content security appliance or other groupware host) as defined in the spam quarantine settings (see the *Cisco Content Security Management Appliance User Guide*). Therefore, regardless of the number of Email Security appliances delivering mail to the Security Management appliance, all released mail, notifications, and alerts are sent to a single host (groupware or content security appliance). Take care not to overburden the primary host for delivery from the Security Management appliance.

## Migrating from a Local Spam Quarantine to an External Quarantine

If you are currently using the local spam quarantine on an Email Security appliance but would like to migrate to an external spam quarantine hosted on a Security Management appliance — while retaining access to the messages in the local quarantine — you should prevent new messages from entering the local quarantine during the transition.

Consider the following possible strategies:

- Configuring anti-spam settings — Configure the anti-spam settings on your mail policy specifying the Security Management appliance as the alternate host. This action sends new spam to the external quarantine while still allowing access to the local quarantine.
- Setting a shorter expiration time — Configure the Schedule Delete After setting on the local quarantine to a shorter duration.
- Deleting all of the remaining messages — To delete all remaining messages in the local quarantine, disable the quarantine and click the “Delete All” link on the local quarantines page (see [Deleting Messages from the Spam Quarantine, page 31-24](#)). This link only becomes available when a local spam quarantine with messages still contained in it has been disabled.

You should now be ready to enable the external quarantine and disable the local quarantine.



**Note**

If both the local quarantine and the external quarantine are enabled, the local quarantine is used.

## Enabling an External Spam Quarantine and External Safelist/Blocklist

You can enable only one external spam quarantine on an Email Security appliance.

**Before You Begin**

- Review the information in [Mail Flow and the External Spam Quarantine, page 42-2](#).
- Review and take action on the information in [Migrating from a Local Spam Quarantine to an External Quarantine, page 42-3](#).
- Configure the Security Management appliance to support the centralized spam quarantine and safelist/blocklist features. See the documentation for your Security Management appliance.
- If a different external spam quarantine was previously configured for the Email Security appliance, first disable the external spam quarantine setting.

Complete the following procedure on each Email Security appliance.

**Procedure**

- 
- Step 1** Select **Security Services > Centralized Services > Spam Quarantine**.
- Step 2** Click **Configure**.
- Step 3** Select **Enable External Spam Quarantine**.
- Step 4** In the Name field, enter the name of the Security Management appliance.  
The name is not significant, and is used for reference only. For example, enter the hostname of the Security Management appliance.
- Step 5** Enter an IP address and port number.  
These must match the IP address and port number that are specified on the Security Management appliance in the Spam Quarantines Settings page (**Management Appliance > Centralized Services > Spam Quarantine**.)
- Step 6** (Optional) Select the check box to enable the **External Safelist/Blocklist** feature, and specify the appropriate blocklist action.
- Step 7** Submit and commit your changes.
- Step 8** Repeat this procedure for each Email Security appliance.
- 

**What To Do Next**

If you have been using a local quarantine, see [Disabling the Local Spam Quarantine to Activate the External Quarantine, page 42-4](#).

**Related Topics**

- [Local Versus External Spam Quarantine, page 31-1](#)
- [Chapter 31, “Spam Quarantine”](#)
- [Chapter 13, “Anti-Spam”](#)
- [How to Configure the Appliance to Scan Messages for Spam, page 13-2](#)

## Disabling the Local Spam Quarantine to Activate the External Quarantine

If you were using a local spam quarantine before enabling an external spam quarantine, you must disable the local quarantine in order to send messages to the external quarantine.

**Before You Begin**

Follow all directions, including information in the Before You Begin section, in [Enabling an External Spam Quarantine and External Safelist/Blocklist](#), page 42-3.

**Procedure**

- 
- |               |                                                                                                                                                                                                            |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Select <b>Monitor &gt; Spam Quarantine</b> .                                                                                                                                                               |
| <b>Step 2</b> | In the Spam Quarantine section, click the <b>Spam Quarantine</b> link.                                                                                                                                     |
| <b>Step 3</b> | Deselect <b>Enable Spam Quarantine</b> .                                                                                                                                                                   |
|               | Ignore any warnings to adjust mail policies as a result of this change. Mail policies automatically send messages to the external spam quarantine if you have configured the external quarantine settings. |
| <b>Step 4</b> | Submit and commit your changes.                                                                                                                                                                            |
- 

## Troubleshooting an External Spam Quarantine

### Email Security Appliance Reprocesses Messages Released from External Quarantine

**Problem** Messages released from the Security Management appliance are unnecessarily reprocessed by the Email Security appliance.

**Solution** This can occur if the IP address of the Security Management appliance has changed. See [Mail Flow and the External Spam Quarantine](#), page 42-2.

## About Centralizing Policy, Virus, and Outbreak Quarantines

- [Centralized Policy, Virus, and Outbreak Quarantines](#), page 42-5
- [About Migration of Policy, Virus, and Outbreak Quarantines](#), page 42-6
- [Centralizing Policy, Virus, and Outbreak Quarantines](#), page 42-7
- [About Disabling Centralized Policy, Virus, and Outbreak Quarantines](#), page 42-8
- [Troubleshooting Centralized Policy, Virus, and Outbreak Quarantines](#), page 42-9

## Centralized Policy, Virus, and Outbreak Quarantines

You can centralize policy, virus, and outbreak quarantines on a Security Management appliance. Messages are processed by Email Security appliances but are stored in quarantines on the Security Management appliance.

Centralizing policy, virus, and outbreak quarantines offers the following benefits:

- Administrators can manage quarantined messages from multiple Email Security appliances in one location.
- Quarantined messages are stored behind the firewall instead of in the DMZ, reducing security risk.

- Centralized quarantines can be backed up using the standard backup functionality on the Security Management appliance.

For complete information, see the user guide or online help for your Security Management appliance.

## Restrictions and Limitations of Centralized Policy, Virus, and Outbreak Quarantines

- On each Email Security appliance, either all policy, virus, and outbreak quarantines must be centralized or all must be stored locally.
- Because scanning engines are not available on Security Management appliances, you cannot manually test messages in policy, virus, or outbreak quarantines for viruses.

## Requirements for Centralized Policy, Virus, and Outbreak Quarantines in Cluster Configurations

You can enable centralized policy, virus, and outbreak quarantines at any level for clustered appliances.

Requirements:

- Before you enable centralized policy, virus, and outbreak quarantines on an Email Security appliance at a particular level (machine, group, or cluster), all appliances that belong to the same level must first be added to the Security Management appliance.
- Content and message filters and DLP message actions must be configured at the same level and not overridden at any level below that level.
- Centralized policy, virus, and outbreak quarantines settings must be configured at the same level and not be overridden at any level below the configured level.
- Ensure that the interface to be used for communications with the Security Management appliance has the same name on all appliances in the group or cluster.

For example:

If you want to enable centralized policy, virus, and outbreak quarantines at the cluster or group level, but an Email Security appliance which is connected to the cluster has these settings defined at the machine level, you must remove the centralized quarantines settings configured at the machine level before you can enable the feature at the cluster or group level.

## About Migration of Policy, Virus, and Outbreak Quarantines

When you centralize policy, virus, and outbreak quarantines, existing policy, virus, and outbreak quarantines on your Email Security appliance migrate to the Security Management appliance.

You will configure migration on the Security Management appliance, but migration occurs when you commit the change enabling centralized policy, virus, and outbreak quarantines on the Email Security appliance.

As soon as you commit this change, the following occur:

- Local policy, virus, and outbreak quarantines on the Email Security appliance are disabled. All new messages entering these quarantines will be quarantined on the Security Management appliance.
- Migration of existing non-spam quarantines to the Security Management appliance begins.
- All local policy, virus, and outbreak quarantines are deleted. If you configured a custom migration, any local policy quarantines that you chose not to migrate are also deleted. For effects of deleting policy quarantines, see [About Deleting Policy Quarantines, page 30-7](#).

- A message that was in multiple quarantines before migration will be in the corresponding centralized quarantines after migration.
- Migration happens in the background. The amount of time it takes depends on the size of your quarantines and on your network. When you enable centralized quarantines on the Email Security appliance, you can enter one or more email addresses that will receive notification when migration is complete.
- The settings in the centralized quarantine, not those of the originating local quarantine, apply to the messages. However, the original expiration time still applies to each message.

**Note**

All centralized quarantines that are automatically created during migration have the default quarantine settings.

## Centralizing Policy, Virus, and Outbreak Quarantines

**Note**

Perform this procedure during a maintenance window or off-peak hours.

### Before You Begin

- You must first configure your Security Management appliance for centralized policy, virus, and outbreak quarantines. See the table in the “Centralizing Policy Virus, and Outbreak Quarantines” section in the “Centralized Policy, Virus, and Outbreak Quarantines” chapter in the online help or user guide for the Security Management appliance.
- If the space allocated to centralized quarantines on the Security Management appliance will be smaller than the amount of space that your existing local quarantines collectively occupy, messages will be expired early based on the quarantine settings on the Security Management appliance. Before migration, consider taking manual action to reduce quarantine sizes. For more information about early expiration, see [Default Actions for Automatically Processed Quarantined Messages](#), page 30-4.
- If you have chosen automatic migration, or configured custom migration to create centralized quarantines during migration, consider noting the current quarantine settings on your Email Security appliances in order to use them as guidelines for configuring the centralized quarantines.
- If your Email Security appliances are deployed in a cluster configuration, see [Requirements for Centralized Policy, Virus, and Outbreak Quarantines in Cluster Configurations](#), page 42-6.
- Be aware of the changes that will occur as soon as you commit the changes in this procedure. See [About Migration of Policy, Virus, and Outbreak Quarantines](#), page 42-6.

### Procedure

- Step 1** Choose **Security Services > Centralized Services > Policy, Virus, and Outbreak Quarantines**.
- Step 2** Click **Enable**.
- Step 3** Enter the interface and port to use for communication with the Security Management appliance. Make sure the interface and port are reachable from the Security Management appliance.  
If your Email Security appliances are clustered, the interface you select must be available on all machines in the cluster.



- Step 4** To receive notification when migration is complete, enter one or more email addresses.
- Step 5** Verify the information about quarantines to be migrated to be sure that this is what you want.
- Step 6** If you are completing a Custom migration, note any quarantines that will be deleted when you commit the changes in this procedure.
- Step 7** Verify that the information about content and message filters and DLP message actions to be updated is as you expect it to be.



**Note** For cluster configurations, filters and message actions can be automatically updated on a particular level only if filters and message actions are defined at that level and not overridden at any level below that level. After migration, you may need to manually reconfigure filters and message actions with centralized quarantine names.

- Step 8** If you need to reconfigure migration mapping:
- Return to the Security Management appliance.
  - Reconfigure the migration mapping.  
On the management appliance, select a quarantines to remap, then click **Remove from Centralized Quarantine**. Then you can remap the quarantine.
  - Commit the new migration configuration on the Security Management appliance.
  - Start this procedure from the beginning.

**Important!** Be sure to reload the **Security Services > Centralized Services > Policy, Virus, and Outbreak Quarantines** page.

- Step 9** Click **Submit**.
- Step 10** If you need to reconfigure migration mapping, follow the procedure in [Step 8](#).
- Step 11** Commit your changes.



**Note** While migration is in progress, avoid making configuration changes on the Email Security appliance or the Security Management appliance.

- Step 12** Look at the top of the page to monitor migration status, or, if you entered an email address when configuring migration, await the email notifying you that migration is complete.

### What To Do Next

Perform the remaining tasks described in the table in the “Centralizing Policy, Virus, and Outbreak Quarantines” topic in the online help or user guide for the Security Management appliance.

### Related Topics

- [Which User Groups Can Access Policy, Virus, and Outbreak Quarantines, page 30-10](#)

## About Disabling Centralized Policy, Virus, and Outbreak Quarantines

When you disable centralized policy, virus, and outbreak quarantines on the Email Security appliance:

- Local quarantines are automatically enabled on the Email Security appliance.

- System-created quarantines and quarantines that are referenced by message filters, content filters, and DLP message actions are automatically created on the Email Security appliance. The Virus, Outbreak, and Unclassified quarantines are created with the same settings that they had before quarantines were centralized, including assigned user roles. All other quarantines are created with default settings.
- Newly quarantined messages go immediately to local quarantines.
- Messages in the centralized quarantine at the time it is disabled remain there until one of the following occurs:
  - Messages are manually deleted or automatically deleted when they expire.
  - Messages are manually or automatically released, if one of the following is also true:
    - \* An alternate release appliance is configured on the Security Management appliance. See the online help or documentation for the Security Management appliance.
    - \* Centralized quarantines are again enabled on the Email Security appliance.

## Disabling Centralized Policy, Virus, and Outbreak Quarantines

### Before You Begin

- Understand the impacts of disabling centralized policy, virus, and outbreak quarantines.
- Do one of the following:
  - Process all messages that are currently in centralized policy, virus, and outbreak quarantines.
  - Ensure that you have designated an alternate release appliance to process messages that are released from the centralized quarantine after you disable it. For information, see the online help or user guide for your Security Management appliance.

### Procedure

- 
- |               |                                                                                                                                           |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | On the Email Security appliance, choose <b>Security Services &gt; Centralized Services &gt; Policy, Virus, and Outbreak Quarantines</b> . |
| <b>Step 2</b> | Disable centralized policy, virus, and outbreak quarantines.                                                                              |
| <b>Step 3</b> | Submit and commit the change.                                                                                                             |
| <b>Step 4</b> | Customize the settings of the newly created local quarantines.                                                                            |
- 

## Troubleshooting Centralized Policy, Virus, and Outbreak Quarantines

### If a Cisco Content Security Management Appliance Goes Out of Service

If Policy, Virus, and Outbreak Quarantines are centralized on a Security Management appliance that goes out of service, you should disable these centralized quarantines on the Email Security appliance.

If you deploy a replacement Security Management appliance, you must reconfigure quarantine migration on the Security Management appliance and on each Email Security appliance. See the table in the “Centralizing Policy Virus, and Outbreak Quarantines” section in the “Centralized Policy, Virus, and Outbreak Quarantines” chapter in the online help or user guide for the Security Management appliance.

# Configuring Centralized Reporting

## Before You Begin

- Enable and configure centralized reporting on a Security Management appliance. See prerequisites and instructions in the *Cisco Content Security Management Appliance User Guide*.
- Ensure that sufficient disk space is allocated to the reporting service on the Security Management appliance.

## Procedure

- 
- |               |                                                                            |
|---------------|----------------------------------------------------------------------------|
| <b>Step 1</b> | Click <b>Security Services &gt; Reporting</b> .                            |
| <b>Step 2</b> | In the Reporting Service section, select the Centralized Reporting option. |
| <b>Step 3</b> | Submit and commit your changes.                                            |
- 

## Requirements for Advanced Malware Protection Reporting

For required configurations for full reporting on Advanced Malware Protection (file reputation and file analysis) features on the Security Management appliance, see the information about Advanced Malware Protection reports in the email reporting chapter of the online help or user guide for your version of the Security Management appliance software.

## Availability of Reporting Information after Changing to Centralized Reporting

When centralized reporting is enabled on an Email Security appliance:

- Existing data on the Email Security appliance for the monthly report is not transferred to the Security Management appliance.
- Archived reports on the Email Security appliance are not available.
- The Email Security appliance stores only a week's worth of data.
- New data for the monthly and yearly reports is stored on the Security Management appliance.
- Scheduled reports on the Email Security appliance are suspended.
- You can no longer access the scheduled report configuration page on the Email Security appliance.

## About Disabling Centralized Reporting

If you disable centralized reporting on the Email Security appliance, the Email Security appliance begins storing new monthly report data, scheduled reports resume, and you can access its archived reports. After disabling centralized reporting, the appliance only displays data for the past hour and day, but not the past week or month. This is temporary. The appliance will display the reports for the past week and month after it accumulates enough data. If the Email Security appliance is placed back into centralized reporting mode, it will display data for the past week in the interactive reports.

# Configuring Centralized Message Tracking

**Note**

You cannot enable both centralized and local tracking on an Email Security appliance.

**Procedure**

- 
- Step 1** Click **Security Services > Message Tracking**.
  - Step 2** In the Message Tracking Service section, click **Edit Settings**.
  - Step 3** Select the **Enable Message Tracking Service** check box.
  - Step 4** Select the Centralized Tracking option.
  - Step 5** (Optional) Select the check box to save information for rejected connections.

**Note**

Saving tracking information for rejected connections can adversely affect the performance of the Security Management appliance.

- Step 6** Submit and commit your changes.
- 

**What To Do Next**

To use centralized tracking, you must enable the feature on the Email Security appliances *and* the Security Management appliance. To enable centralized tracking on the Security Management appliance, see the *Cisco Content Security Management Appliance User Guide*.

## Using Centralized Services

For instructions on using centralized services, see the *Cisco Content Security Management Appliance User Guide*.





## FTP, SSH, SCP, and Telnet Access

You can access any interface you create on the appliance through a variety of services.

- [IP Interfaces, page A-1](#)
- [Configuring FTP Access to the Email Security Appliance, page A-2](#)
- [Secure Copy \(scp\) Access, page A-5](#)
- [Accessing the Email Security appliance via a Serial Connection, page A-5](#)

## 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 assign an Internet Protocol version 4 (IPv4) or version 6 (IPv6) to an IP interface or both.

**Table A-1** *Services Enabled by Default on Interfaces*

| Service | Default port | Enabled by default?               |                           |
|---------|--------------|-----------------------------------|---------------------------|
|         |              | Management interface <sup>a</sup> | New interfaces you create |
| FTP     | 21           | No                                | No                        |
| Telnet  | 23           | Yes                               | No                        |
| SSH     | 22           | Yes                               | No                        |
| HTTP    | 80           | Yes                               | No                        |
| HTTPS   | 443          | Yes                               | No                        |

a.The “Management Interface” settings shown here are also the default settings for the Data 1 Interface on Cisco C170 appliances.

- If you need to access the appliance via the graphical user interface (GUI), you must enable HTTP and/or HTTPS on an interface.
- If you need to access the appliance for the purposes of uploading or downloading configuration files, you must enable FTP or Telnet on an interface.
- You can also upload or download files using secure copy (`scp`).

You can configure HTTP or HTTPS 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. You can also “join” interfaces into distinct groups (via the CLI), and the system will cycle through these groups when delivering email.

Joining or grouping Virtual Gateways is useful for load-balancing large email campaigns across several interfaces. You can also create VLANs, and configure them just as you would any other interface (via the CLI). For more information, see [Chapter 37, “Advanced Network Configuration.”](#)

- 

## How AsyncOS Selects Default IP Interface

AsyncOS selects the default IP interface based on the order in which the IP interfaces appear under **Network > IP Interfaces** page or in the `ifconfig` CLI command. The first IP interface in the list that resides on the subnet in question is used.

If there are multiple IP addresses configured within the same subnet as the default gateway, the IP address with the lowest number is used. For example, if the following IP addresses are configured within the same subnet,

- 10.10.10.2/24
- 10.10.10.30/24
- 10.10.10.100/24
- 10.10.10.105/24

AsyncOS chooses 10.10.10.2/24 as the default IP interface.

## Configuring FTP Access to the Email Security Appliance

### Procedure

- 
- Step 1** Use the Network > IP Interfaces page or the `interfaceconfig` command to enable FTP access for the interface.

**WARNING:** By disabling services via the `interfaceconfig` command, you have the potential to disconnect yourself from the CLI, depending on how you are connected to the appliance. Do not disable services with this command if you are not able to reconnect to the appliance using another protocol, the Serial interface, or the default settings on the Management port.

In this example, the Management interface is edited to enable FTP access on port 21 (the default port):

**Figure A-1** *Edit IP Interface Page*  
**Edit IP Interface**

| IP Interface Settings |                                            |      |
|-----------------------|--------------------------------------------|------|
| Name:                 | Management                                 |      |
| Ethernet Port:        | Management                                 |      |
| IP Address:           | 172.19.0.11 *                              |      |
| Netmask:              | 255.255.255.0 *                            |      |
| Hostname:             | elroy.run                                  |      |
| Services:             | Service                                    | Port |
|                       | <input checked="" type="checkbox"/> FTP    | 21   |
|                       | <input checked="" type="checkbox"/> Telnet | 23   |
|                       | <input checked="" type="checkbox"/> SSH    | 22 * |



**Note** Remember to `commit` your changes before moving on to the next step.

**Step 2** Access the interface via FTP. Ensure you are using the correct IP address for the interface. For example:

```
$ ftp 192.168.42.42
```



**Note** Many browsers also allow you to access interfaces via FTP.

**Step 3** Browse to the directory for the specific task you are trying to accomplish. After you have accessed an interface via FTP, you can browse the following directories to copy and add (“GET” and “PUT”) files. See the following table.



| Directory Name                                                                                                                                                                                                                                                                                                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| /configuration                                                                                                                                                                                                                                                                                                       | <p>The directory where data from the following commands is exported to and/or imported (saved) from:</p> <ul style="list-style-type: none"> <li>Virtual Gateway mappings (<code>altsrchost</code>)</li> <li>configuration data in XML format (<code>saveconfig</code>, <code>loadconfig</code>)</li> <li>Host Access Table (HAT) (<code>hostaccess</code>)</li> <li>Recipient Access Table (RAT) (<code>rcptaccess</code>)</li> <li>SMTP routes entries (<code>smtproutes</code>)</li> <li>alias tables (<code>aliasconfig</code>)</li> <li>masquerading tables (<code>masquerade</code>)</li> <li>message <b>filters</b> (<code>filters</code>)</li> <li>global unsubscribe data (<code>unsubscribe</code>)</li> <li>test messages for the <code>trace</code> command</li> <li>Safelist/Blocklist backup file, saved in the following format: <code>slbl&lt;timestamp&gt;&lt;serial number&gt;.csv</code></li> </ul> |
| /antivirus                                                                                                                                                                                                                                                                                                           | The directory where the Anti-Virus engine log files are kept. You can inspect the log files this directory to manually check for the last successful download of the virus definition file ( <code>scan.dat</code> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| /configuration<br>/system_logs<br>/cli_logs<br>/status<br>/reportd_logs<br>reportqueryd_logs<br>/ftpd_logs<br>/mail_logs<br>/asarchive<br>/bounces<br>/error_logs<br>/avarchive<br>/gui_logs<br>/sntpd_logs<br>/RAID.output<br>/euq_logs<br>/scanning<br>/antispam<br>/antivirus<br>/euqgui_logs<br>/ipmitool.output | <p>Created automatically for <b>logging</b> via the <code>logconfig</code> and <code>rollovernow</code> commands. See <a href="#">Logging</a> for a detailed description of each log.</p> <p>See “Log File Type Comparison” for the differences between each log file type.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

**Step 4** Use your FTP program to upload and download files to and from the appropriate directory.

## Secure Copy (scp) Access

If your client operating system supports a secure copy (`scp`) command, you can copy files to and from the directories listed in the previous table. For example, in the following example, the file `/tmp/test.txt` is copied from the client machine to the configuration directory of the appliance with the hostname of `mail3.example.com`.

Note that the command prompts for the password for the user (`admin`). This example is shown for reference only; your particular operating system's implementation of secure copy may vary.

```
% scp /tmp/test.txt admin@mail3.example.com:configuration
```

```
The authenticity of host 'mail3.example.com (192.168.42.42)' can't be established.
```

```
DSA key fingerprint is 69:02:01:1d:9b:eb:eb:80:0c:a1:f5:a6:61:da:c8:db.
```

```
Are you sure you want to continue connecting (yes/no)? yes
```

```
Warning: Permanently added 'mail3.example.com ' (DSA) to the list of known hosts.
```

```
admin@mail3.example.com's password: (type the password)
```

```
test.txt 100% |*****| 1007 00:00
```

```
%
```

In this example, the same file is copied from the appliance to the client machine:

```
% scp admin@mail3.example.com:configuration/text.txt .
```

```
admin@mail3.example.com's password: (type the password)
```

```
test.txt 100% |*****| 1007 00:00
```

```
%
```

You can use secure copy (`scp`) as an alternative to FTP to transfer files to and from the Cisco appliance.



### Note


Only users in the operators and administrators group can use secure copy (`scp`) to access the appliance. For more information, see [Adding Users, page 32-4](#).

## Accessing the Email Security appliance via a Serial Connection

If you are connecting to the appliance via a serial connection, use the following information for the console port.

Complete information about this port is in the hardware installation guide for your appliance.

## Pinout Details for the Serial Port in 80- Series Hardware



## Pinout Details for the Serial Port in 70-Series Hardware

Figure A-2 illustrates the pin numbers for the serial port connector, and Table A-2 defines the pin assignments and interface signals for the serial port connector.




**Table A-2** Serial Port Pin Assignments

| Pin   | Signal | I/O | Definition          |
|-------|--------|-----|---------------------|
| 1     | DCD    | I   | Data carrier detect |
| 2     | SIN    | I   | Serial input        |
| 3     | SOUT   | O   | Serial output       |
| 4     | DTR    | O   | Data terminal ready |
| 5     | GND    | n/a | Signal ground       |
| 6     | DSR    | I   | Data set ready      |
| 7     | RTS    | I   | Request to send     |
| 8     | CTS    | O   | Clear to send       |
| 9     | RI     | I   | Ring indicator      |
| Shell | n/a    | n/a | Chassis ground      |



# Assigning Network and IP Addresses

This appendix describes general rules on networks and IP address assignments, and it presents some strategies for connecting the Cisco appliance to your network.

- [Ethernet Interfaces, page B-1](#)
- [Selecting IP Addresses and Netmasks, page B-1](#)
- [Strategies for Connecting Your Cisco Appliance, page B-3](#)

## Ethernet Interfaces

The Cisco X1070, C670, and C370 appliances are equipped with as many as four Ethernet interfaces located on the rear panel of the system depending on the configuration (whether or not you have the optional optical network interface). They are labeled:

- Management
- Data1
- Data2
- Data3
- Data4

The Cisco C170 appliance is equipped with two Ethernet interfaces located on the rear panel of the system. They are labeled:

- Data1
- Data2

## Selecting IP Addresses and Netmasks

When you configure the network, the Cisco appliance must be able to uniquely select an interface to send an outgoing packet. This requirement will drive some of the decisions regarding IP address and netmask selection for the Ethernet interfaces. The rule is that only one interface can be on a single network (as determined through the applications of netmasks to the IP addresses of the interfaces).

An IP address identifies a physical interface on any given network. A physical Ethernet interface can have more than one IP address for which it accepts packets. An Ethernet interface that has more than one IP address can send packets over that interface with any one of the IP addresses as the source address in the packet. This property is used in implementing Virtual Gateway technology.

The purpose of a netmask is to divide an IP address into a network address and a host address. The network address can be thought of as the network part (the bits matching the netmask) of the IP address. The host address is the remaining bits of the IP address. The number of bits in a four octet address that are significant are sometimes expressed in CIDR (Classless Inter-Domain Routing) style. This is a slash followed by the number of bits (1-32).

A netmask can be expressed in this way by simply counting the ones in binary, so 255.255.255.0 becomes “/24” and 255.255.240.0 becomes “/20”.

## Sample Interface Configurations

This section shows sample interface configurations based on some typical networks. The example will use two interfaces called Int1 and Int2. In the case of the Cisco appliance, these interface names can represent any two interfaces out of the three Cisco interfaces (Management, Data1, Data2).

### Network 1:

Separate interfaces must appear to be on separate networks.

| Interface | IP address   | netmask       | net address    |
|-----------|--------------|---------------|----------------|
| Int1      | 192.168.1.10 | 255.255.255.0 | 192.168.1.0/24 |
| Int2      | 192.168.0.10 | 255.255.255.0 | 192.168.0.0/24 |

Data addressed to 192.168.1.x (where X is any number 1-255, except for your own address, 10 in this case) will go out on Int1. Anything addressed to 192.168.0.x will go out on Int2. Any packet headed for some other address not in these formats, most likely out on a WAN or the Internet, will be sent to the default gateway which must itself be on one of these networks. The default gateway will then forward the packet on.

### Network 2:

The network addresses (network parts of the IP addresses) of two different interfaces cannot be the same.

| Ethernet Interface | IP address   | netmask     | net address    |
|--------------------|--------------|-------------|----------------|
| Int1               | 192.168.1.10 | 255.255.0.0 | 192.168.0.0/16 |
| Int2               | 192.168.0.10 | 255.255.0.0 | 192.168.0.0/16 |

This situation presents a conflict in that two different Ethernet interfaces have the same network address. If a packet from the Cisco appliance is sent to 192.168.1.11, there is no way to decide which Ethernet interface should be used to deliver the packet. If the two Ethernet interfaces are connected to two separate physical networks, the packet may be delivered to the incorrect network and never find its destination. The Cisco appliance will not allow you to configure your network with conflicts.

You can connect two Ethernet interfaces to the same physical network, but you must construct IP addresses and netmasks to allow the Cisco appliance to select a unique delivery interface.

## IP Addresses, Interfaces, and Routing

When selecting an interface on which to perform a command or function in the GUI or CLI that allows you to select an interface (for example, upgrading AsyncOS, or configuring DNS, etc.), routing (your default gateway) will take precedence over your selection.

For example, suppose you have an Cisco appliance with the 3 network interfaces configured, each on a different network segment (assume all /24):

| Ethernet   | IP           |
|------------|--------------|
| Management | 192.19.0.100 |
| data1      | 192.19.1.100 |
| data2      | 192.19.2.100 |

And your Default gateway is 192.19.0.1.

Now, if you perform an AsyncOS upgrade (or other command or function that allows you to select an interface) and you select the IP that is on data1 (192.19.1.100), you would expect all the TCP traffic to occur over the data1 ethernet interface. However, what happens is that the traffic will go out of the interface that is set as your default gateway, in this case Management, but will be stamped with the source address of the IP on data1.

## Summary

The Cisco appliance must always be able to identify a unique interface over which a packet will be delivered. To make this decision, the Cisco appliance uses a combination of the packet's destination IP address, and the network and IP address settings of its Ethernet interfaces. The following table summarizes the preceding examples:

|                              | Same Network | Different Network |
|------------------------------|--------------|-------------------|
| Same Physical Interface      | Allowed      | Allowed           |
| Different Physical Interface | Not Allowed  | Allowed           |

## Strategies for Connecting Your Cisco Appliance

Keep these things in mind when connecting your Cisco appliance:

- Administrative traffic (CLI, web interface, log delivery) traffic is usually small compared to email traffic.
- If two Ethernet interfaces are connected to the same network switch, but end up talking to a single interface on another host downstream, or are connected to a network hub where all data are echoed to all ports, no advantage is gained by using two interfaces.
- SMTP conversations over an interface operating at 1000Base-T will be slightly faster than ones over the same interfaces operating at 100Base-T, but only under ideal conditions.
- There is no point in optimizing connections to your network if there is a bottleneck in some other part of your delivery network. Bottlenecks most often occur in the connection to the Internet and further upstream at your connectivity provider.

The number of Cisco appliance interfaces that you choose to connect and how you address them should be dictated by the complexity of your underlying network. It is not necessary to connect multiple interfaces if your network topology or data volumes do not call for it. It is also possible to keep the connection simple at first as you familiarize yourself with the gateway and then increase the connectivity as volume and network topology require it.



## Example of Mail Policies and Content Filters

### Overview of Incoming Mail Policies

The following example demonstrates the features of mail policies by illustrating the following tasks:

1. Editing the anti-spam, anti-virus, Outbreak Filter, and Content Filters for the default Incoming Mail Policy.
2. Adding two new policies for different sets of users — the sales organization and the engineering organization — and then configuring different email security settings for each.
3. Creating three new content filters to be used in the Incoming Mail Overview policy table.
4. Editing the policies again to enable the content filters for some groups, but not for others.

This example is meant to show the power and flexibility with which you can manage different recipient-based settings for anti-spam, anti-virus, Outbreak Filter, and Content Filters for mail policies. This example assigns these a custom user role called “Policy Administrator” that has mail policy and content filters access privileges. For more detailed information about how anti-spam, anti-virus, Outbreak filters, and delegated administration work, refer to the chapters following this one:

- [Anti-Spam, page 13-1](#)
- [Anti-Virus, page 12-1](#)
- [Outbreak Filters, page 14-1](#)
- [Distributing Administrative Tasks, page 32-1](#)

### Accessing Mail Policies

You can access incoming and outgoing mail policies by using the Mail Policies menu.

On brand new systems, if you completed all steps in the system setup wizard and you chose to enable Anti-Spam, Sophos or McAfee Anti-Virus, and Outbreak Filters, the Incoming Mail Policies Page will resemble [Figure C-1](#).

By default, these settings are enabled for the default Incoming Mail Policy:

- Anti-Spam (if the Spam Quarantine is enabled): Enabled
  - Positively-identified spam: quarantine, prepend the message subject
  - Suspected spam: quarantine, prepend the message subject
  - Marketing email: scanning not enabled
- Anti-Spam (if the Spam Quarantine is not enabled): Enabled



- Positively-identified spam: deliver, prepend the message subject
- Suspected spam: deliver, prepend the message subject
- Marketing email: scanning not enabled
- Anti-Virus: Enabled, Scan and Repair viruses, include an X-header with anti-virus scanning results
  - Repaired messages: deliver, prepend the message subject
  - Encrypted messages: deliver, prepend the message subject
  - Unscannable messages: deliver, prepend the message subject
  - Virus infected messages: drop
- Outbreak Filters: Enabled
  - No file extensions are excepted
  - Retention time for messages with suspect viral attachments is 1 day
  - Message modification is not enabled
- Content Filters: Disable

**Figure C-1** Incoming Mail Policies Page: Defaults for a Brand New Appliance  
Incoming Mail Policies

| Find Policies  |  |  |                                                                 |               |  |  |
|----------------|--|--|-----------------------------------------------------------------|---------------|--|--|
| Email Address: |  |  | <input type="radio"/> Recipient<br><input type="radio"/> Sender | Find Policies |  |  |

| Policies      |                |                                                                                                     |                                                                              |                 |                                 |        |
|---------------|----------------|-----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|-----------------|---------------------------------|--------|
| Add Policy... |                |                                                                                                     |                                                                              |                 |                                 |        |
| Order         | Policy Name    | Anti-Spam                                                                                           | Anti-Virus                                                                   | Content Filters | Outbreak Filters                | Delete |
|               | Default Policy | IronPort Anti-Spam<br>Positive: Quarantine<br>Suspected: Quarantine<br>Marketing Messages: Disabled | Sophos<br>Encrypted: Deliver<br>Unscannable: Deliver<br>Virus Positive: Drop | Disabled        | Retention Time:<br>Virus: 1 day |        |

Key: Default Custom Readonly



#### Note

In this example, the Incoming Mail Policy will use the default anti-spam settings for when the Spam Quarantine is enabled.

## Enabled, Disabled, and “Not Available”

The columns in a mail policy table (either incoming or outgoing) display links for the state of the security service for each policy name. If a service is enabled, the word “Enabled” or a summary of the configuration is displayed. Similarly, the word “Disabled” is displayed if a service is disabled.

“Not Available” is displayed as a link if the license agreement for a service has not been accepted yet or a service has expired. In these cases, clicking the “Not Available” link will display the global page within the Security Services tab, rather than the page where you can configure per-policy settings for a service. An alert is displayed to let you know that your page has changed to a different tab. See [Figure C-2](#).

**Figure C-2 Security Services Not Available**  
Incoming Mail Policies

**Find Policies**

Email Address:  Recipient ☒ Sender ☐ Find Policies

---

**Policies**

Add Policy...

| Order | Policy Name    | Anti-Spam     | Anti-Virus    | Content Filters | Outbreak Filters | Delete |
|-------|----------------|---------------|---------------|-----------------|------------------|--------|
|       | Default Policy | Not Available | Not Available | Disabled        | Not Available    |        |

Key: Default Custom Readonly

## Configuring the Default Anti-Spam Policies for Incoming Messages

Each row in the mail policy table represents a different policy. Each column represents a different security service.

- To edit the default policy, click any of the links for a security service in the bottom row of the incoming or outgoing mail policy table.

In this example, you will change the anti-spam settings for the default policy for incoming mail to be more aggressive. The default value is to quarantine positively identified and suspected spam messages, with marketing email scanning disabled. This example shows how to change the setting so that positively identified spam is dropped. Suspected spam continues to be quarantined. Marketing email scanning is enabled, with marketing messages being delivered to the intended recipients. The subjects of marketing messages will be prepended with the text `[MARKETING]`.

### Procedure

**Step 1** Click the link for the anti-spam security service.



**Note** For default security service settings, the first setting on the page defines whether the service is enabled for the policy. You can click “Disable” to disable the service altogether.

**Step 2** In the “Positively Identified Spam Settings” section, change the “Action to apply to this message” to Drop.

**Step 3** In the “Marketing Email Settings” section, click **Yes** to enable marketing email scanning.

If enabled, the default action is to deliver legitimate marketing messages while prepending the subject with the text `[MARKETING]`.

The “Add text to message” field only accepts US-ASCII characters.

**Step 4** Click **Submit**. Note that the summary link for the anti-spam security service in the Incoming Mail Policies table has changed to reflect the new values.

Similar to the steps above, you can change the default anti-virus and virus outbreak filter settings for the default policy.

**Figure C-3 Anti-Spam Settings Page**  
Mail Policies: Anti-Spam

| Anti-Spam Settings                                                                          |                                                                                                           |
|---------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| <b>Policy:</b>                                                                              | Default                                                                                                   |
| Enable Anti-Spam Scanning for This Policy:                                                  | <input checked="" type="radio"/> Use IronPort Anti-Spam service<br><input type="radio"/> Disabled         |
| <b>Positively-Identified Spam Settings</b>                                                  |                                                                                                           |
| Apply This Action to Message:                                                               | Drop                                                                                                      |
| Add Text to Subject:                                                                        | Prepend [SPAM]                                                                                            |
| Advanced Optional settings for custom header and message delivery.                          |                                                                                                           |
| <b>Suspected Spam Settings</b>                                                              |                                                                                                           |
| Enable Suspected Spam Scanning:                                                             | <input type="radio"/> No <input checked="" type="radio"/> Yes                                             |
| Apply This Action to Message:                                                               | Spam Quarantine                                                                                           |
| Note: If local and external quarantines are defined, mail will be sent to local quarantine. |                                                                                                           |
| Add Text to Subject:                                                                        | Prepend [SUSPECTED SPAM]                                                                                  |
| Advanced Optional settings for custom header and message delivery.                          |                                                                                                           |
| <b>Marketing Email Settings</b>                                                             |                                                                                                           |
| Enable Marketing Email Scanning:                                                            | <input type="radio"/> No <input checked="" type="radio"/> Yes                                             |
| Apply This Action to Message:                                                               | Deliver                                                                                                   |
| Send to Alternate Host (optional):                                                          |                                                                                                           |
| Add Text to Subject:                                                                        | Prepend [MARKETING]                                                                                       |
| Advanced Optional settings for custom header and message delivery.                          |                                                                                                           |
| <b>Spam Thresholds</b>                                                                      |                                                                                                           |
| Spam is scored on a 1-100 scale. The higher the score, the more likely a message is a spam. |                                                                                                           |
| IronPort Anti-Spam:                                                                         | <input checked="" type="radio"/> Use the Default Thresholds<br><input type="radio"/> Use Custom Settings: |
| Positively Identified Spam:                                                                 | Score > 90 (50 - 100)                                                                                     |
| Suspected Spam:                                                                             | Score > 50 (minimum 25, cannot exceed positive spam score)                                                |

Cancel Submit

## Creating a Mail Policy for a Group of Sender and Recipients

In this part of the example, you will create two new policies: one for the sales organization (whose members will be defined by an LDAP acceptance query), and another for the engineering organization. Both policies will be assigned to the Policy Administrator custom user role to make delegated administrators belonging to this role responsible for managing these policies. You will then configure different email security settings for each.

### Procedure

- Step 1** Click the **Add Policy** button to begin creating a new policy.
- Step 2** Define a unique name for and adjust the order of the policy (if necessary).
 

The name of the policy must be unique to the Mail Policies table (either incoming or outgoing) in which it is defined.

Remember that each recipient is evaluated for each policy in the appropriate table (incoming or outgoing) in a top-down fashion.
- Step 3** Click the Editable by (Roles) link and select the custom user roles for the delegated administrators who will be responsible for managing the mail policy.
 

When you click the link, AsyncOS displays the custom roles for delegated administrators that have edit privileges for mail policies. Delegated administrators can edit a policy's Anti-Spam, Anti-Virus, and Outbreak Filters settings and enable or disable content filters for the policy. Only operators and administrators can modify a mail policy's name or its senders, recipients, or groups. Custom user roles that have full access to mail policies are automatically assigned to mail policies.

See the [Distributing Administrative Tasks](#) for more information on delegated administration.

**Step 4** Define users for the policy.

You define whether the user is a sender or a recipient. (See [Examples of Policy Matching](#), page 10-4 for more detail.) The form shown in [Figure C-4](#) defaults to recipients for incoming mail policies and to senders for outgoing mail policies.

Users for a given policy can be defined in the following ways:

- Full email address: `user@example.com`
- Partial email address: `user@`
- All users in a domain: `@example.com`
- All users in a partial domain: `@.example.com`
- By matching an LDAP Query

**Note**

Entries for users are case-insensitive in both the GUI and CLI in AsyncOS. For example, if you enter the recipient `Joe@` for a user, a message sent to `joe@example.com` will match.

If you store user information within LDAP directories in your network infrastructure — for example, in Microsoft Active Directory, SunONE Directory Server (formerly known as “iPlanet Directory Server”), or Open LDAP directories — you can configure the 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.

See [LDAP Queries](#) for more information.

**Figure C-4**      **Defining Users for a Policy**  
**Add Incoming Mail Policy**

**Step 5** Click the **Add** button to add users into the Current Users list.

Policies can contain mixtures of senders, recipients, and LDAP queries.

Use the **Remove** button to remove a defined user from the list of current users.

**Step 6** When you are finished adding users, click **Submit**.

Note that all security services settings are set to use the default values when you first add a policy.

**Figure C-5** Newly Added Policy – Sales Group

| Policies      |                |                                                                                                     |                                                                              |                 |                                 |        |
|---------------|----------------|-----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|-----------------|---------------------------------|--------|
| Add Policy... |                |                                                                                                     |                                                                              |                 |                                 |        |
| Order         | Policy Name    | Anti-Spam                                                                                           | Anti-Virus                                                                   | Content Filters | Outbreak Filters                | Delete |
| 1             | Sales_Team     | (use default)                                                                                       | (use default)                                                                | (use default)   | (use default)                   |        |
|               | Default Policy | IronPort Anti-Spam<br>Positive: Quarantine<br>Suspected: Quarantine<br>Marketing Messages: Disabled | Sophos<br>Encrypted: Deliver<br>Unscannable: Deliver<br>Virus Positive: Drop | Disabled        | Retention Time:<br>Virus: 1 day |        |

**Step 7** Click the **Add Policy** button again to add another new policy.

In this policy, individual email addresses for members of the engineering team are defined:

**Figure C-6** Creating a Policy for the Engineering Team  
Add Incoming Mail Policy

Policy Name: ?

Engineering  
(e.g. my IT policy)

Editable by (Roles):

Policy Administrator

Insert Before Policy:

2 (Default Policy)

Sender

Recipient ?

Email Address(es)

bob@example.com

mary@example.com

fred@example.com

(e.g. user@example.com, user@, @example.com, @example.com)

LDAP Group Query

Query: Sales\_West.group

Group:

Add >

Remove

Current Users

Recipient: bob@example.com

Recipient: mary@example.com

Recipient: fred@example.com

Cancel

Submit

**Step 8** When you are finished adding users for the engineering policy, click **Submit**.

**Step 9** Commit your changes.

**Figure C-7** Newly Added Policy — Engineering Team

| Policies      |                |                                                                                                     |                                                                              |                 |                                 |        |
|---------------|----------------|-----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|-----------------|---------------------------------|--------|
| Add Policy... |                |                                                                                                     |                                                                              |                 |                                 |        |
| Order         | Policy Name    | Anti-Spam                                                                                           | Anti-Virus                                                                   | Content Filters | Outbreak Filters                | Delete |
| 1             | Sales_Team     | (use default)                                                                                       | (use default)                                                                | (use default)   | (use default)                   |        |
| 2             | Engineering    | (use default)                                                                                       | (use default)                                                                | (use default)   | (use default)                   |        |
|               | Default Policy | IronPort Anti-Spam<br>Positive: Quarantine<br>Suspected: Quarantine<br>Marketing Messages: Disabled | Sophos<br>Encrypted: Deliver<br>Unscannable: Deliver<br>Virus Positive: Drop | Disabled        | Retention Time:<br>Virus: 1 day |        |

**Note**

At this point, both newly created policies have the same settings applied to them as those in the default policy. Messages to users of either policy will match; however, the mail processing settings are not any different from the default policy. Therefore, messages that match users in the “Sales\_Group” or “Engineering” policies will not be processed any differently than the default policy.

## Default, Custom, and Disabled

The key at the bottom of the table shows how the color coding of cells for specific policies relates to the policy defined for the default row:

Key: Default Custom Disabled

- Yellow shading shows that the policy is using the same settings as the default policy.
- No shading (white) shows that the policy is using different settings than the default policy.
- Grey shading shows that the security service has been disabled for the policy.

## Creating Mail Policies for Different Groups of Senders and Recipients

In this part of the example, you will edit the two policies just created in the previous section.

- For the sales group, you will change the anti-spam settings to be even more aggressive than the default policy. (See [Configuring the Default Anti-Spam Policies for Incoming Messages, page C-3](#).) The default policy of dropping positively identified spam messages will be kept. However, in this example, you will change the setting for marketing messages so that they will be sent to the Spam quarantine.

This aggressive policy has the effect of minimizing unwanted messages being sent to sales team inboxes.

See [Anti-Spam, page 13-1](#) for more information on anti-spam settings.

- For the engineering team, customize the Outbreak Filters feature setting so that it will modify the URLs in suspicious messages, except for links to example.com. Attachment files with the extension “dwg” will be bypassed by the Outbreak Filter scanning.

See [Outbreak Filters, page 14-1](#) for more information on configuring Outbreak Filters.

To edit the anti-spam settings for the sales team policy:

### Procedure

- Step 1** Click the link for the Anti-Spam security service (the Anti-Spam) column in the sales policy row.

Because the policy was just added, the link is named: (use default).

**Figure C-8** *Editing the Anti-Spam Settings for the Sales Team Policy*

| Policies      |                |                                                                                              |
|---------------|----------------|----------------------------------------------------------------------------------------------|
| Add Policy... |                |                                                                                              |
| Order         | Policy Name    | Anti-Spam                                                                                    |
| 1             | Sales_Team     | (use default)                                                                                |
| 2             | Engineering    | (use default)                                                                                |
|               | Default Policy | IronPort Anti-Spam<br>Positive: Drop<br>Suspected: Quarantine<br>Marketing Messages: Deliver |

- Step 2** On the anti-spam security service page, change the value for “Enable Anti-Spam Scanning for this Policy” from “Use Default Settings” to “Use Anti-Spam service.”

Choosing “Use Anti-Spam service” here allows you to override the settings defined in the default policy.

- Step 3** In the “Positively-Identified Spam Settings” section, change the “Apply This Action to Message” to “Drop.”

- Step 4** In the “Suspected Spam Settings” section, click **Yes** to enable suspected spam scanning.

- Step 5** In the “Suspected Spam Settings” section, change the “Apply This Action to Message” to “Spam Quarantine.”



**Note** Selecting the Spam quarantine forwards mail according to the settings defined in the Spam Quarantine chapter.

- Step 6** In the “Add text to subject” field, click **None**.

Messages delivered to the Spam quarantine will have no additional subject tagging.

- Step 7** In the “Marketing Email Settings” section, click **Yes** to enable scanning for marketing mail from legitimate sources.

- Step 8** In the “Apply This Action to Message” section, select “Spam Quarantine.”

- Step 9** Submit and commit your changes.

Not that the shading shows that the policy is using different settings than the default policy.

**Figure C-9** *Anti-Spam Settings for the Sales Group Policy Changed*

| Policies      |                |                                                                                                 |
|---------------|----------------|-------------------------------------------------------------------------------------------------|
| Add Policy... |                |                                                                                                 |
| Order         | Policy Name    | Anti-Spam                                                                                       |
| 1             | Sales_Team     | IronPort Anti-Spam<br>Positive: Drop<br>Suspected: Quarantine<br>Marketing Messages: Quarantine |
| 2             | Engineering    | (use default)                                                                                   |
|               | Default Policy | IronPort Anti-Spam<br>Positive: Drop<br>Suspected: Quarantine<br>Marketing Messages: Deliver    |

At this point, any message that is suspected spam and whose recipient matches the LDAP query defined for the sales team policy will be delivered to the Spam Quarantine.

To edit the Outbreak Filter settings for the engineering team policy:

## Procedure

- Step 1** Click the link for the Outbreak Filters feature security service (the Outbreak Filters column) in the engineering policy row.

Because the policy was just added, the link is named: (use default).

**Figure C-10** Editing the Outbreak Filters Feature Settings for the Engineering Team Policy

| Policies      |                |                                                                                                     |                                                                              |                 |                                 |        |
|---------------|----------------|-----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|-----------------|---------------------------------|--------|
| Add Policy... |                |                                                                                                     |                                                                              |                 |                                 |        |
| Order         | Policy Name    | Anti-Spam                                                                                           | Anti-Virus                                                                   | Content Filters | Outbreak Filters                | Delete |
| 1             | Sales_Team     | IronPort Anti-Spam<br>Positive: Drop<br>Suspected: Quarantine<br>Marketing Messages: Quarantine     | (use default)                                                                | (use default)   | (use default)                   |        |
| 2             | Engineering    | (use default)                                                                                       | (use default)                                                                | (use default)   | (use default)                   |        |
|               | Default Policy | IronPort Anti-Spam<br>Positive: Quarantine<br>Suspected: Quarantine<br>Marketing Messages: Disabled | Sophos<br>Encrypted: Deliver<br>Unscannable: Deliver<br>Virus Positive: Drop | Disabled        | Retention Time:<br>Virus: 1 day |        |

- Step 2** On the Outbreak Filters feature security service page, change the scanning setting for the policy to “Enable Outbreak Filtering (Customize settings).”

Choosing “(Customize settings)” here allows you to override the settings defined in the default policy.

Doing so will also enable the contents of the rest of the page to allow you to select different settings.

- Step 3** In the “Bypass Attachment Scanning” section of the page, type **dwg** in the in the file extension field.

The file extension “**dwg**” is not in the list of known file type that the appliance can recognize by its fingerprint when attachment scanning.



**Note** You do not need to type the period (.) before the three letter filename extension.

- Step 4** Click **Add Extension** to add .dwg files to the list of file extensions that will bypass Outbreak Filters feature scanning.

- Step 5** Click **Enable Message Modification**.

Enabling message modification allows the appliance to scan for targeted threats, such as phishing and scams, and URLs to suspicious or malicious websites. The appliance can rewrite links in messages to redirect the user through the Cisco Security proxy if they attempt to access the website.



**Note** Anti-spamming scanning must be enabled on the mail policy in order for Outbreak Filters to scan for targeted, non-viral threats.

- Step 6** Select for **Enable for Unsigned Messages**.

This allows the appliance to rewrite URLs in signed messages. You must enable URL rewriting to be able to configure other Message Modification settings and the length of time that messages found to be non-viral threats stay in the quarantine before being released. This example uses the default retention time of 4 hours.

- Step 7** Enter example.com in the **Bypass Domain Scanning** field.

The appliance will not modify links to example.com.

- Step 8** Select System Generated for the **Threat Disclaimer**.

The appliance can insert a disclaimer above the message body to warn the user about the message’s contents. This example uses the system generated threat disclaimer.



**Figure C-11 Outbreak Filters Settings**  
Mail Policies: Outbreak Filters

Outbreak Filtering for Policy: Sales\_Team

Enable Outbreak Filtering (Customize settings)

---

**Outbreak Filter Settings**

Quarantine Threat Level: 3

Maximum Quarantine Retention: Viral Attachments: 1 Days  
Other Threats: 4 Hours

Bypass Attachment Scanning: ☒ Select File Extension... ☐   
Add Extension

File Extensions to Bypass: None defined

---

**Message Modification**

☒ Enable Message Modification

Message Modification Threat Level: 3

Message Subject: Prepend [MODIFIED FOR PROTECTION]

URL Rewriting: Cisco Security proxy scans and rewrites suspicious or malicious URLs.  
☒ Enable only for unsigned messages (recommended)  
☐ Enable for all messages  
☐ Disable

Bypass Domain Scanning: example.com  
(examples: example.com, crm.example.com, 10.0.0.1, 10.0.0.0/24)

Threat Disclaimer: System Generated  
Preview Disclaimer

Disclaimer text will be applied to the top of the message body for Suspicious and Quarantined messages. To create custom disclaimers go to Mail Policies > Text Resources

Cancel Submit

**Step 9** Submit and commit your changes.

Note that the shading shows that the policy is using different settings than the default policy.

**Figure C-12 Virus Filters Settings for the Engineering Policy Changed**

| Policies |                |                                                                                                     |                                                                              |                 |                                                   |
|----------|----------------|-----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|-----------------|---------------------------------------------------|
| Order    | Policy Name    | Anti-Spam                                                                                           | Anti-Virus                                                                   | Content Filters | Outbreak Filters                                  |
| 1        | Sales_Team     | IronPort Anti-Spam<br>Positive: Drop<br>Suspected: Quarantine<br>Marketing Messages: Quarantine     | (use default)                                                                | (use default)   | (use default)                                     |
| 2        | Engineering    | (use default)                                                                                       | (use default)                                                                | (use default)   | Retention Time:<br>Virus: 1 day<br>Other: 4 hours |
|          | Default Policy | IronPort Anti-Spam<br>Positive: Quarantine<br>Suspected: Quarantine<br>Marketing Messages: Disabled | Sophos<br>Encrypted: Deliver<br>Unscannable: Deliver<br>Virus Positive: Drop | Disabled        | Retention Time:<br>Virus: 1 day                   |

At this point, any message that contains an attachment whose file extension is `dwg` — and whose recipient matches the recipients defined for the engineering team policy — will bypass the Outbreak Filter scanning and continue processing. Messages that contain links to the `example.com` domain will not have their links modified to redirect through the Cisco Security proxy and will not be considered suspicious.

## Finding Senders or Recipients in Mail Policies

Use the “Find Policies” button to search for users already defined in policies defined in the Incoming or Outgoing Mail Policies pages.

For example, typing `joe@example.com` and clicking the Find Policies button will display results showing which policies contain defined users that will match the policy.

**Figure C-13 Finding Users in Policies**

**Find Policies**

Email Address:  Recipient ☒ Sender ☐ Find Policies

Results: Email Address "Recipient: bob@example.com" is defined in the following policies:

- Engineering
- Default Policy (all users)

**Policies matching "bob@example.com"**

Add Policy... Show All Policies

| Order | Policy Name    | Anti-Spam                                                                                            | Anti-Virus                                                                   | Content Filters | Outbreak Filters                                  | Delete |
|-------|----------------|------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|-----------------|---------------------------------------------------|--------|
| 2     | Engineering    | (use default)                                                                                        | (use default)                                                                | (use default)   | Retention Time:<br>Virus: 1 day<br>Other: 4 hours |        |
|       | Default Policy | IronPort: Anti-Spam<br>Positive: Quarantine<br>Suspected: Quarantine<br>Marketing Messages: Disabled | Sophos<br>Encrypted: Deliver<br>Unscannable: Deliver<br>Virus Positive: Drop | Disabled        | Retention Time:<br>Virus: 1 day                   |        |

Click the name of the policy to jump to the Edit Policy page to edit the users for that policy.

Note that the default policy will always be shown when you search for any user, because, by definition, if a sender or recipient does not match any other configured policies, it will *always* match the default policy.

## Managed Exceptions

Using the steps shown in the two examples above, you can begin to create and configure policies on a *managed exception* basis. In other words, after evaluating your organization's needs you can configure policies so that the majority of messages will be handled by the default policy. You can then create additional "exception" policies for specific users or user groups, managing the differing policies as needed. In this manner, message splintering will be minimized and you are less likely to impact system performance from the processing of each splinter message in the work queue.

You can define policies based on your organizations' or users' tolerance for spam, viruses, and policy enforcement. [Table C-1 on page C-11](#) outlines several example policies. "Aggressive" policies are designed to minimize the amount of spam and viruses that reach end-users mailboxes. "Conservative" policies are tailored to avoid false positives and prevent users from missing messages, regardless of policies.

**Table C-1 Aggressive and Conservative Mail Policy Settings**

|                      | Aggressive Settings                                                                                                                         | Conservative Settings                                                                                                                                   |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Anti-Spam</b>     | Positively identified spam: Drop<br>Suspected spam: Quarantine<br>Marketing mail: Deliver and prepend "[Marketing]" to the subject messages | Positively identified spam: Quarantine<br>Suspected spam: Deliver and prepend "[Suspected Spam]" to the subject of messages<br>Marketing mail: Disabled |
| <b>Anti-Virus</b>    | Repaired messages: Deliver<br>Encrypted messages: Drop<br>Unscannable messages: Drop<br>Infectious messages: Drop                           | Repaired messages: Deliver<br>Encrypted messages: Quarantine<br>Unscannable messages: Quarantine<br>Infectious messages: Drop                           |
| <b>Virus Filters</b> | Enabled, no specific filename extensions or domains allowed to bypass<br>Enable message modification for all messages                       | Enabled with specific filename extensions or domains allowed to bypass<br>Enable message modification for unsigned messages                             |

## Filtering Messages Based on Content

In this part of the example, you will create three new content filters to be used in the Incoming Mail Policy table. All of these content filters will be editable by delegated administrators belonging to the Policy Administration custom user role. You will create the following:

1. “scan\_for\_confidential”

This filter will scan messages for the string “confidential.” If the string is found, a copy of the message will be sent to email alias `hr@example.com`, and the message will be sent to the Policy quarantine area.

2. “no\_mp3s”

This filter will strip MP3 attachments and notify the recipients that an MP3 file was stripped.

3. “ex\_employee”

This content filter will scan for messages sent to a specific envelope recipient address (an ex-employee). If the message matches, a specific notification message will be sent to the sender of the message and then the message will be bounced.

After creating the content filters, you will then configure each of the policies (including the default policy) to enable the specific content filters in differing combinations.

## Quarantining Message with “Confidential” in the Subject

The first example content filter contains one condition and two actions.

### Procedure

- 
- Step 1** Click the Mail Policies tab.
  - Step 2** Click Incoming Content Filters.
  - Step 3** Click the **Add Filter** button.
  - Step 4** In the Name field, type `scan_for_confidential` as the name of the new filter.  
Filter names can contain ASCII characters, numbers, underscores or dashes. The first character of a content filter name must be a letter or an underscore.
  - Step 5** Click the **Editable By (Roles)** link, select the Policy Administrator and click **OK**.  
Delegated administrators who belong to the Policy Administrator user role will be able to edit this content filter and use it in their mail policies.
  - Step 6** In the Description field, type the description. For example: `scan all incoming mail for the string 'confidential'.`
  - Step 7** Click **Add Condition**.
  - Step 8** Select Message Body.
  - Step 9** Type `confidential` in the Contains text: field and click **OK**.  
The Add Content Filter page shows the condition added.
  - Step 10** Click **Add Action**.
  - Step 11** Select Send Copy To (Bcc:).
  - Step 12** In the Email Addresses field, type `hr@example.com`.

**Step 13** In the Subject field, type `[message matched confidential filter]`.

**Step 14** Click **OK**.

The Add Content Filter page shows the action added.

**Step 15** Click **Add Action**.

**Step 16** Select Quarantine.

**Step 17** In the drop-down menu, select the Policy quarantine area.

**Step 18** Click **OK**.

The Add Content Filter page shows the second action added.

**Step 19** Submit and commit your changes.

At this point, the content filter is not enabled for any incoming Mail Policy; in this example, you have only added a new content filter to the master list. Because it has not been applied to any policy, no email processing by the appliance will be affected by this filter.

## Stripping MP3 Attachments from Messages

The second example content filter contains no conditions and one action.

### Procedure

**Step 1** Click the **Add Filter** button.

**Step 2** In the Name field, type `no_mp3s` as the name of the new filter.

**Step 3** Click the **Editable By (Roles)** link, select the Policy Administrator and click **OK**.

**Step 4** In the Description field, type the description. For example: `strip all MP3 attachments`.

**Step 5** Click **Add Action**.

**Step 6** Select Strip Attachment by File Info.

**Step 7** Select `File type is`.

**Step 8** In the drop-down field, select `-- mp3`.

**Step 9** Enter a replacement message if desired.

**Step 10** Click **OK**.

**Step 11** Submit and commit your changes.



### Note

It is not necessary to specify a condition when creating a content filter. When no condition is defined, any actions defined will always apply in the rule. (Specifying no condition is equivalent to using the `true()` message filter rule — all messages will be matched if the content filter is applied to a policy.)

## Bouncing Messages Sent to a Former Employee

The third content filter example uses one condition and two actions.

### Procedure

- 
- Step 1** Click the **Add Filter** button.
- Step 2** In the Name: field, type `ex_employee` as the name of the new filter.
- Step 3** Click the **Editable By (Roles)** link, select the Policy Administrator and click **OK**.
- Step 4** In the Description: field, type the description. For example: `bounce messages intended for Doug`.
- Step 5** Click **Add Condition**.
- Step 6** Select Envelope Recipient.
- Step 7** For the envelope recipient, select `Begins with`, and type `doug@`.
- Step 8** Click **OK**.
- The Content Filters page refreshes to show the condition added. Note that you could create an LDAP directory containing the email addresses of former employees. As ex-employees are added to that directory, this content filter would be dynamically updated.
- Step 9** Click **Add Action**.
- Step 10** Select Notify.
- Step 11** Select the checkbox for Sender and, in the Subject field, type `message bounced for ex-employee of example.com`.
- Step 12** In the Use template section, select a notification template.
- 
- Note** Some sections of the content filter rule builder will not appear in the user interface if the resource has not been preconfigured. For example, content dictionaries, notification templates, and message disclaimers will not appear as options if they have not been configured previously via the Mail Policies > Dictionaries page (or the `dictionaryconfig` command in the CLI). For more information about creating dictionaries, see [Content Dictionaries, page 21-2](#).
- 
- Step 13** Click **OK**.
- The Add Content Filters page shows the action added.
- Step 14** Click **Add Action**.
- Step 15** Select Bounce (Final Action) and click **OK**.
- You can only specify one final action for a content filter. If you try to attempt to add more than one final action, the GUI displays an error.
- Adding this action may will cause senders of messages to this ex-employee to potentially receive two messages: one for the notification template, and one for the bounce notification template.
- Step 16** Submit and commit your changes.

## Applying Individual Content Filters to Different Groups of Recipients

In the examples above, you created three content filters using the Incoming Content Filters pages. The Incoming Content Filters and Outgoing Content filters pages hold the “master lists” of all possible content filters that can be applied to a policy.

**Figure C-14 Incoming Content Filters: Three Filters Created**  
**Incoming Content Filters**

| Filters       |                       |                                                      |           |        |  |
|---------------|-----------------------|------------------------------------------------------|-----------|--------|--|
| Add Filter... |                       |                                                      |           |        |  |
| Order         | Filter Name           | Description   Rules   Policies                       | Duplicate | Delete |  |
| 1             | scan_for_confidential | scan all incoming mail for the string 'confidential' |           |        |  |
| 2             | no_mp3s               | strip all MP3 attachments                            |           |        |  |
| 3             | ex_employee           | bounce messages intended for Doug                    |           |        |  |

In this part of the example, you will apply the three new content filters to be used in the Incoming Mail Policy table.

- The default policy will receive all three content filters.
- The engineering group will *not* receive the no\_mp3s filter.
- The sales group will receive the content filters as the default incoming mail policy.

## Enabling Content Filters for All Recipients by Default

Click the links to enable and select content filters for individual policies.

### Procedure

- Step 1** Click Incoming Mail Policies to return to the Incoming Mail Policy table.
- The page is refreshed to show the default policy and the two policies added in [Creating a Mail Policy for a Group of Sender and Recipients, page C-4](#). Note that content filtering is disabled by default for all policies.
- Step 2** Click the link for the Content Filters security service (the Content Filters column) in the default policy row. See [Figure C-15](#).

**Figure C-15 Editing the Content Filters Setting for the Default Incoming Mail Policy**

| Policies      |                |                                                                                                     |                                                                              |                 |                                                   |        |
|---------------|----------------|-----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|-----------------|---------------------------------------------------|--------|
| Add Policy... |                |                                                                                                     |                                                                              |                 |                                                   |        |
| Order         | Policy Name    | Anti-Spam                                                                                           | Anti-Virus                                                                   | Content Filters | Outbreak Filters                                  | Delete |
| 1             | Sales_Team     | IronPort Anti-Spam<br>Positive: Drop<br>Suspected: Quarantine<br>Marketing Messages: Quarantine     | (use default)                                                                | (use default)   | (use default)                                     |        |
| 2             | Engineering    | (use default)                                                                                       | (use default)                                                                | (use default)   | Retention Time:<br>Virus: 1 day<br>Other: 4 hours |        |
|               | Default Policy | IronPort Anti-Spam<br>Positive: Quarantine<br>Suspected: Quarantine<br>Marketing Messages: Disabled | Sophos<br>Encrypted: Deliver<br>Unscannable: Deliver<br>Virus Positive: Drop | Disabled        | Retention Time:<br>Virus: 1 day                   |        |

- Step 3** On the Content Filtering security service page, change the value Content Filtering for Default Policy from “Disable Content Filters” to “Enable Content Filters (Customize settings).”

**Figure C-16** Enabling Content Filters for the Policy and Selecting Specific Content Filters  
Mail Policies: Content Filters

Content Filtering for: Default Policy

Enable Content Filters (Customize settings) ☒   
 Enable Content Filters (Customize settings) ☐   
 Disable Content Filters ☐

| Order | Filter Name           | Description                                          | Enable                   |
|-------|-----------------------|------------------------------------------------------|--------------------------|
| 1     | scan_for_confidential | scan all incoming mail for the string 'confidential' | <input type="checkbox"/> |
| 2     | no_mp3s               | strip all MP3 attachments                            | <input type="checkbox"/> |
| 3     | ex_employee           | bounce messages intended for Doug                    | <input type="checkbox"/> |

Cancel Submit

The content filters defined in the master list (which were created in [Overview of Content Filters, page 11-1](#) using the Incoming Content Filters pages) are displayed on this page. When you change the value to “Enable Content Filters (Customize settings),” the checkboxes for each filter change from disabled (greyed out) to become enabled.

**Step 4** Check the **Enable** checkbox for each content filter.

**Step 5** Click **Submit**.

The table on the Incoming Mail Policies page shows the names of the filters that have been enabled for the default policy.

**Figure C-17** Three Content Filters Enabled for the Default Incoming Mail Policy

|                |                                                                                                     |                                                                              |                                                 |
|----------------|-----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|-------------------------------------------------|
| Default Policy | IronPort Anti-Spam<br>Positive: Quarantine<br>Suspected: Quarantine<br>Marketing Messages: Disabled | Sophos<br>Encrypted: Deliver<br>Unscannable: Deliver<br>Virus Positive: Drop | scan_for_confidential<br>no_mp3s<br>ex_employee |
|----------------|-----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|-------------------------------------------------|

## Allowing MP3 Attachments for Recipients in Engineering

To disable the “no\_mp3s” content filters for the “engineering” policy:

### Procedure

**Step 1** Click the link for the Content Filters security service (the Content Filters column) in the engineering team policy row.

**Step 2** On the Content Filtering security service page, change the value for Content Filtering for Policy: Engineering from “Enable Content Filtering (Inherit default policy settings)” to “Enable Content Filtering (Customize settings).”

Because this policy was using the default values, when you change the value from “Use Default Settings” to “Yes,” the checkboxes for each filter change from disabled (greyed out) to become enabled.

**Step 3** Deselect the checkbox for the “no\_mp3s” filter.

**Figure C-18** Deselecting a Content Filter  
Mail Policies: Content Filters

Content Filtering for Policy: Engineering

Enable Content Filters (Customize settings) ☒

| Order | Filter Name           | Description                                          | Enable                              |
|-------|-----------------------|------------------------------------------------------|-------------------------------------|
| 1     | scan_for_confidential | scan all incoming mail for the string 'confidential' | <input checked="" type="checkbox"/> |
| 2     | no_mp3s               | strip all MP3 attachments                            | <input type="checkbox"/>            |
| 3     | ex_employee           | bounce messages intended for Doug                    | <input checked="" type="checkbox"/> |

Cancel Submit

**Step 4** Click **Submit**.

The table on the Incoming Mail Policies page shows the names of the filters that have been enabled for the engineering policy.

**Figure C-19 Updated Content Filters for Incoming Mail Policies**

| Policies      |                |                                                                                                     |                                                                              |                                                 |                                                   |        |
|---------------|----------------|-----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|-------------------------------------------------|---------------------------------------------------|--------|
| Add Policy... |                |                                                                                                     |                                                                              |                                                 |                                                   |        |
| Order         | Policy Name    | Anti-Spam                                                                                           | Anti-Virus                                                                   | Content Filters                                 | Outbreak Filters                                  | Delete |
| 1             | Sales_Team     | IronPort Anti-Spam<br>Positive: Drop<br>Suspected: Quarantine<br>Marketing Messages: Quarantine     | (use default)                                                                | (use default)                                   | (use default)                                     |        |
| 2             | Engineering    | (use default)                                                                                       | (use default)                                                                | scan_for_confidential<br>ex_employee            | Retention Time:<br>Virus: 1 day<br>Other: 4 hours |        |
|               | Default Policy | IronPort Anti-Spam<br>Positive: Quarantine<br>Suspected: Quarantine<br>Marketing Messages: Disabled | Sophos<br>Encrypted: Deliver<br>Unscannable: Deliver<br>Virus Positive: Drop | scan_for_confidential<br>no_mp3s<br>ex_employee | Retention Time:<br>Virus: 1 day                   |        |

**Step 5** Commit your changes.

At this point, incoming messages that match the user list for the engineering policy will not have MP3 attachments stripped; however, all other incoming messages will have MP3 attachments stripped.

## Notes on Configuring Content Filters in the GUI

- It is not necessary to specify a condition when creating a content filter. When no action is defined, any actions defined will always apply in the rule. (Specifying no action is equivalent to using the `true()` message filter rule — all messages will be matched if the content filter is applied to a policy.)
- If you do not assign a custom user role to a content filter, the content filter is public and can be used by any delegated administrator for their mail policies. See [Distributing Administrative Tasks](#) for more information on delegated administrators and content filters.
- Administrators and operators can view and edit all content filters on an appliance, even when the content filters are assigned to custom user roles.
- When entering text for filter rules and actions, the following meta characters have special meaning in regular expression matching: `. ^ $ * + ? { [ ] \ | ( )`

If you do not wish to use regular expression you should use a `\` (backslash) to escape any of these characters. For example: `"\*Warning\*"`

- When you define more than one Condition for a content filter, you can define whether *all* of the defined actions (that is, a logical AND) or any of the defined actions (logical OR) need to apply in order for the content filter to be considered a match.



**Figure C-20** Choosing Any or All of the Following Conditions

| Add Filter                  |                                                                                                                        |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------|
| Name:                       | <input type="text"/>                                                                                                   |
| Currently used by policies: |                                                                                                                        |
| Description:                | <input type="text"/>                                                                                                   |
| Order:                      | 5 <input type="button" value="v"/>                                                                                     |
| Apply filter:               | <input checked="" type="radio"/> If one or more conditions match<br><input type="radio"/> Only if ALL conditions match |

- You can test message splintering and content filters by creating “benign” content filters. For example, it is possible to create a content filter whose only action is “deliver.” This content filter will not affect mail processing; however, you can use this filter to test how the mail policy processing affects other elements in the system (for example, the mail logs).
- Conversely, using the “master list” concept of the Incoming or Outgoing Content Filters, it is possible to create very powerful, wide-sweeping content filters that will immediately affect message processing for all mail handled by the appliance. The process for this is to:
  - Use the Incoming or Outgoing Content Filters page to create a new content filter whose order is 1.
  - Use the Incoming or Outgoing Mail Policies page to enable the new content filter for the default policy.
  - Enable the content filter for all remaining policies.
- The Bcc: and Quarantine actions available in Content Filters can help you determine the retention settings of quarantines you create. (See [Chapter 30, “Policy, Virus, and Outbreak Quarantines.”](#)) You can create filters that would simulate mail flow into and out of your policy quarantines so that messages are not released too quickly from the system (that is, the quarantine areas do not fill their allotted disk space too quickly).
- Because it uses the same settings as the Scan Behavior page or the `scanconfig` command, the “Entire Message” condition does not scan a message’s headers; choosing the “Entire Message” will scan only the message body and attachments. Use the “Subject” or “Header” conditions to search for specific header information.
- Configuring users by LDAP query will only appear in the GUI if you have LDAP servers configured on the appliance (that is, you have configured the appliance to query specific LDAP servers with specific strings using the `ldapconfig` command).
- Some sections of the content filter rule builder will not appear in the GUI if the resource has not been preconfigured. For example, notification templates and message disclaimers will not appear as options if they have not been configured previously using the Text Resources page or the `textconfig` command in the CLI.
- Content filters features will recognize, can contain, and/or scan for text in the following character encodings:
  - Unicode (UTF-8)
  - Unicode (UTF-16)
  - Western European/Latin-1 (ISO 8859-1)
  - Western European/Latin-1 (Windows CP1252)
  - Traditional Chinese (Big 5)
  - Simplified Chinese (GB 2312)
  - Simplified Chinese (HZ GB 2312)

- Korean (ISO 2022-KR)
- Korean (KS-C-5601/EUC-KR)
- Japanese (Shift-JIS (X0123))
- Japanese (ISO-2022-JP)
- Japanese (EUC)

You can mix and match multiple character sets within a single content filter. Refer to your web browser's documentation for help displaying and entering text in multiple character encodings. Most browsers can render multiple character sets simultaneously.

**Figure C-21 Multiple Character Sets in a Content Filter**

The screenshot shows a web interface for configuring content filters. Under the 'Conditions' tab, there is a dropdown menu set to 'Entire Message', a 'Contains' dropdown, and a text input field containing 'Hello, 你好吗? My name'. Below the input field is an 'Add Condition' button. A list of conditions is shown below, with the first one highlighted: 'body-contains("Hello, 你好吗? My name is Steve")'.

- On the Incoming or Outgoing Content Filters summary pages, use the links for “Description,” “Rules,” and “Policies” to change the view presented for the content filters:
  - The **Description** view shows the text you entered in the description field for each content filter. (This is the default view.)
  - The **Rules** view shows the rules and regular expressions build by the rule builder page.
  - The **Policies** shows the policies for which each content filter is enabled.

**Figure C-22 Using the Links to Toggle Description, Rules, and Policy for Content Filters Incoming Content Filters**

| Filters       |                        |                                                                                                                                                       |           |        |  |
|---------------|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|--------|--|
| Add Filter... |                        |                                                                                                                                                       |           |        |  |
| Order         | Filter Name            | Description   Rules   Policies                                                                                                                        | Duplicate | Delete |  |
| 1             | scan_for_confidential  | scan_for_confidential: if (body-contains("confidential")) { quarantine ("Policy"); bcc ("hr@example.com", "[message matched confidential filter]"); } |           |        |  |
| 2             | no_mp3s                | no_mp3s: if (true) { drop-attachments-by-filetype("mp3", "mp3 deleted"); }                                                                            |           |        |  |
| 3             | ex_employee            | ex_employee: if (rcpt-to == "^doug@") { notify-copy ("{\$EnvelopeSender", "message bounced for ex-employee of example.com"); bounce(); }              |           |        |  |
| 4             | drop_large_attachments | drop_large_attachments: if (true) { drop-attachments-by-size(5242880, "This attachment was too big!"); }                                              |           |        |  |





## Firewall Information

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

**Table D-1** *Firewall Ports*

| Port  | Protocol | In/Out    | Hostname                | Description                                                                                                                                                                                                               |
|-------|----------|-----------|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 20/21 | TCP      | In or Out | AsyncOS IPs, FTP Server | FTP for aggregation of log files.<br>Data ports TCP 1024 and higher must also all be open.<br>For more information, search for FTP port information in the Knowledge Base. See <a href="#">Knowledge Base, page 1-9</a> . |
| 22    | TCP      | In        | AsyncOS IPs             | SSH access to the CLI, aggregation of log files.                                                                                                                                                                          |
| 22    | TCP      | Out       | SSH Server              | SSH aggregation of log files.                                                                                                                                                                                             |
| 22    | TCP      | Out       | SCP Server              | SCP Push to log server                                                                                                                                                                                                    |
| 23    | Telnet   | In        | AsyncOS IPs             | Telnet access to the CLI, aggregation of log files.                                                                                                                                                                       |
| 23    | Telnet   | Out       | Telnet Server           | Telnet upgrades, aggregation of log files (not recommended).                                                                                                                                                              |
| 25    | TCP      | Out       | Any                     | SMTP to send email.                                                                                                                                                                                                       |
| 25    | TCP      | In        | AsyncOS IPs             | SMTP to receive bounced email or if injecting email from outside firewall.                                                                                                                                                |
| 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.                                                                                                    |
| 80    | HTTP     | In        | AsyncOS IPs             | HTTP access to the GUI for system monitoring.                                                                                                                                                                             |
| 80    | HTTP     | Out       | downloads.ironport.com  | Service updates, except for AsyncOS upgrades and McAfee definitions.                                                                                                                                                      |
| 80    | HTTP     | Out       | updates.ironport.com    | AsyncOS upgrades and McAfee Anti-Virus definitions.                                                                                                                                                                       |

**Table D-1** *Firewall Ports (continued)*

|             |       |          |                                                                                                                                                            |                                                                                                                                                                        |
|-------------|-------|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 80          | HTTP  | Out      | cdn-microudates.cloudmark.com                                                                                                                              | Used for updates to third-party spam component in Intelligent MultiScan. Appliance must also connect to CIDR range 208.83.136.0/22 for third-party phone home updates. |
| 82          | HTTP  | In       | AsyncOS IPs                                                                                                                                                | Used for viewing the Cisco Anti-Spam quarantine.                                                                                                                       |
| 83          | HTTPS | In       | AsyncOS IPs                                                                                                                                                | Used for viewing the Cisco Anti-Spam quarantine.                                                                                                                       |
| 110         | TCP   | Out      | POP Server                                                                                                                                                 | POP authentication for end users for Cisco Spam Quarantine                                                                                                             |
| 123         | UDP   | In & Out | NTP Server                                                                                                                                                 | NTP if time servers are outside firewall.                                                                                                                              |
| 143         | TCP   | Out      | IMAP Server                                                                                                                                                | IMAP authentication for end users for Cisco Spam Quarantine                                                                                                            |
| 161         | UDP   | In       | AsyncOS IPs                                                                                                                                                | SNMP Queries                                                                                                                                                           |
| 162         | UDP   | Out      | Management Station                                                                                                                                         | SNMP Traps                                                                                                                                                             |
| 389<br>3268 | LDAP  | Out      | LDAP Servers                                                                                                                                               | LDAP if LDAP directory servers are outside firewall. LDAP authentication for Cisco Spam Quarantine                                                                     |
| 636<br>3269 | LDAPS | Out      | LDAPS                                                                                                                                                      | LDAPS — ActiveDirectory's Global Catalog Server (uses SSL)                                                                                                             |
| 443         | TCP   | In       | AsyncOS IPs                                                                                                                                                | Secure HTTP (https) access to the GUI for system monitoring.                                                                                                           |
| 443         | TCP   | Out      | res.cisco.com                                                                                                                                              | Cisco Registered Envelope Service                                                                                                                                      |
| 443         | TCP   | Out      | update-manifests.ironport.com                                                                                                                              | Verify the latest files for the update server.                                                                                                                         |
| 443         | TCP   | Out      | phonehome.senderbase.org                                                                                                                                   | Receive/Send Outbreak Filters                                                                                                                                          |
| 443         | TCP   | Out      | In the command-line interface, run the <code>websecurityadvancedconfig</code> command and accept all defaults. The Web security service hostname is shown. | Cloud service for obtaining URL reputation and category information for URL filtering.                                                                                 |
| 443         | TCP   | Out      | As configured in Security Services > File Reputation and Analysis, Advanced section, Cloud Server Pool parameter.                                          | If configured, the port for access to cloud services for obtaining file reputation.<br><br>The default port is 32137.<br><br>For file analysis services, see port 443. |

**Table D-1** *Firewall Ports (continued)*

|                 |         |          |                                                                                                                   |                                                                                                                                                                          |
|-----------------|---------|----------|-------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 443             | TCP     | Out      | As configured in Security Services > File Reputation and Analysis, Advanced section.                              | Access to cloud services for file analysis.<br>For file reputation services, see port 443 or 32137.                                                                      |
| 514             | UDP/TCP | Out      | Syslog Server                                                                                                     | Syslog logging                                                                                                                                                           |
| 628             | TCP     | In       | AsyncOS IPs                                                                                                       | QMQP if injecting email from outside firewall.                                                                                                                           |
| 1024 and higher | —       | —        | —                                                                                                                 | See information above for Port 21 (FTP.)                                                                                                                                 |
| 2222            | CCS     | In & Out | AsyncOS IPs                                                                                                       | Cluster Communication Service (for Centralized Management).                                                                                                              |
| 6025            | TCP     | Out      | AsyncOS IPs                                                                                                       | Cisco Spam Quarantine                                                                                                                                                    |
| 7025            | TCP     | In & Out | AsyncOS IPs                                                                                                       | Pass policy, virus, and outbreak quarantine data between Email Security appliances and the Cisco Content Security Management appliance when this feature is centralized. |
| 32137           | TCP     | Out      | As configured in Security Services > File Reputation and Analysis, Advanced section, Cloud Server Pool parameter. | Default port for access to cloud services for obtaining file reputation. See also port 443.<br>For file analysis services, see port 443.                                 |





## End User License Agreement

---

### Cisco Systems End User License Agreement

**IMPORTANT: PLEASE READ THIS END USER LICENSE AGREEMENT CAREFULLY. IT IS VERY IMPORTANT THAT YOU CHECK THAT YOU ARE PURCHASING CISCO SOFTWARE OR EQUIPMENT FROM AN APPROVED SOURCE AND THAT YOU, OR THE ENTITY YOU REPRESENT (COLLECTIVELY, THE "CUSTOMER") HAVE BEEN REGISTERED AS THE END USER FOR THE PURPOSES OF THIS CISCO END USER LICENSE AGREEMENT. IF YOU ARE NOT REGISTERED AS THE END USER YOU HAVE NO LICENSE TO USE THE SOFTWARE AND THE LIMITED WARRANTY IN THIS END USER LICENSE AGREEMENT DOES NOT APPLY. ASSUMING YOU HAVE PURCHASED FROM AN APPROVED SOURCE, DOWNLOADING, INSTALLING OR USING CISCO OR CISCO-SUPPLIED SOFTWARE CONSTITUTES ACCEPTANCE OF THIS AGREEMENT.**

CISCO SYSTEMS, INC. OR ITS SUBSIDIARY LICENSING THE SOFTWARE INSTEAD OF CISCO SYSTEMS, INC. ("CISCO") IS WILLING TO LICENSE THIS SOFTWARE TO YOU ONLY UPON THE CONDITION THAT YOU PURCHASED THE SOFTWARE FROM AN APPROVED SOURCE AND THAT YOU ACCEPT ALL OF THE TERMS CONTAINED IN THIS END USER LICENSE AGREEMENT PLUS ANY ADDITIONAL LIMITATIONS ON THE LICENSE SET FORTH IN A SUPPLEMENTAL LICENSE AGREEMENT ACCOMPANYING THE PRODUCT OR AVAILABLE AT THE TIME OF YOUR ORDER (COLLECTIVELY THE "AGREEMENT"). TO THE EXTENT OF ANY CONFLICT BETWEEN THE TERMS OF THIS END USER LICENSE AGREEMENT AND ANY SUPPLEMENTAL LICENSE AGREEMENT, THE SUPPLEMENTAL LICENSE AGREEMENT SHALL APPLY. BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE, YOU ARE REPRESENTING THAT YOU PURCHASED THE SOFTWARE FROM AN APPROVED SOURCE AND BINDING YOURSELF TO THE AGREEMENT. IF YOU DO NOT AGREE TO ALL OF THE TERMS OF THE AGREEMENT, THEN CISCO IS UNWILLING TO LICENSE THE SOFTWARE TO YOU AND (A) YOU MAY NOT DOWNLOAD, INSTALL OR USE THE SOFTWARE, AND (B) YOU MAY RETURN THE SOFTWARE (INCLUDING ANY UNOPENED CD PACKAGE AND ANY WRITTEN MATERIALS) FOR A FULL REFUND, OR, IF THE SOFTWARE AND WRITTEN MATERIALS ARE SUPPLIED AS PART OF ANOTHER PRODUCT, YOU MAY RETURN THE ENTIRE PRODUCT FOR A FULL REFUND. YOUR RIGHT TO RETURN AND REFUND EXPIRES 30 DAYS AFTER PURCHASE FROM AN APPROVED SOURCE, AND APPLIES ONLY IF YOU ARE THE ORIGINAL AND REGISTERED END USER PURCHASER. FOR THE PURPOSES OF THIS END USER LICENSE AGREEMENT, AN "APPROVED SOURCE" MEANS (A) CISCO; OR (B) A DISTRIBUTOR OR SYSTEMS INTEGRATOR AUTHORIZED BY CISCO TO DISTRIBUTE / SELL CISCO EQUIPMENT, SOFTWARE AND SERVICES WITHIN YOUR TERRITORY TO END USERS; OR (C) A RESELLER AUTHORIZED BY ANY SUCH DISTRIBUTOR OR SYSTEMS



INTEGRATOR IN ACCORDANCE WITH THE TERMS OF THE DISTRIBUTOR'S AGREEMENT WITH CISCO TO DISTRIBUTE / SELL THE CISCO EQUIPMENT, SOFTWARE AND SERVICES WITHIN YOUR TERRITORY TO END USERS.

*THE FOLLOWING TERMS OF THE AGREEMENT GOVERN CUSTOMER'S USE OF THE SOFTWARE (DEFINED BELOW), EXCEPT TO THE EXTENT: (A) THERE IS A SEPARATE SIGNED CONTRACT BETWEEN CUSTOMER AND CISCO GOVERNING CUSTOMER'S USE OF THE SOFTWARE, OR (B) THE SOFTWARE INCLUDES A SEPARATE "CLICK-ACCEPT" LICENSE AGREEMENT OR THIRD PARTY LICENSE AGREEMENT AS PART OF THE INSTALLATION OR DOWNLOAD PROCESS GOVERNING CUSTOMER'S USE OF THE SOFTWARE. TO THE EXTENT OF A CONFLICT BETWEEN THE PROVISIONS OF THE FOREGOING DOCUMENTS, THE ORDER OF PRECEDENCE SHALL BE (1) THE SIGNED CONTRACT, (2) THE CLICK-ACCEPT AGREEMENT OR THIRD PARTY LICENSE AGREEMENT, AND (3) THE AGREEMENT. FOR PURPOSES OF THE AGREEMENT, "SOFTWARE" SHALL MEAN COMPUTER PROGRAMS, INCLUDING FIRMWARE AND COMPUTER PROGRAMS EMBEDDED IN CISCO EQUIPMENT, AS PROVIDED TO CUSTOMER BY AN APPROVED SOURCE, AND ANY UPGRADES, UPDATES, BUG FIXES OR MODIFIED VERSIONS THERETO (COLLECTIVELY, "UPGRADES"), ANY OF THE SAME WHICH HAS BEEN RELICENSED UNDER THE CISCO SOFTWARE TRANSFER AND RE-LICENSING POLICY (AS MAY BE AMENDED BY CISCO FROM TIME TO TIME) OR BACKUP COPIES OF ANY OF THE FOREGOING.*

**License.** Conditioned upon compliance with the terms and conditions of the Agreement, Cisco grants to Customer a nonexclusive and nontransferable license to use for Customer's internal business purposes the Software and the Documentation for which Customer has paid the required license fees to an Approved Source. "Documentation" means written information (whether contained in user or technical manuals, training materials, specifications or otherwise) pertaining to the Software and made available by an Approved Source with the Software in any manner (including on CD-Rom, or on-line). In order to use the Software, Customer may be required to input a registration number or product authorization key and register Customer's copy of the Software online at Cisco's website to obtain the necessary license key or license file.

Customer's license to use the Software shall be limited to, and Customer shall not use the Software in excess of, a single hardware chassis or card or such other limitations as are set forth in the applicable Supplemental License Agreement or in the applicable purchase order which has been accepted by an Approved Source and for which Customer has paid to an Approved Source the required license fee (the "Purchase Order").

Unless otherwise expressly provided in the Documentation or any applicable Supplemental License Agreement, Customer shall use the Software solely as embedded in, for execution on, or (where the applicable Documentation permits installation on non-Cisco equipment) for communication with Cisco equipment owned or leased by Customer and used for Customer's internal business purposes. No other licenses are granted by implication, estoppel or otherwise.

For evaluation or beta copies for which Cisco does not charge a license fee, the above requirement to pay license fees does not apply.

**General Limitations.** This is a license, not a transfer of title, to the Software and Documentation, and Cisco retains ownership of all copies of the Software and Documentation. Customer acknowledges that the Software and Documentation contain trade secrets of Cisco or its suppliers or licensors, including but not limited to the specific internal design and structure of individual programs and associated interface information. Except as otherwise expressly provided under the Agreement, Customer shall only use the Software in connection with the use of Cisco equipment purchased by the Customer from an Approved Source and Customer shall have no right, and Customer specifically agrees not to:

- (i) transfer, assign or sublicense its license rights to any other person or entity (other than in compliance with any Cisco relicensing/transfer policy then in force), or use the Software on Cisco equipment not purchased by the Customer from an Approved Source or on secondhand Cisco equipment, and Customer acknowledges that any attempted transfer, assignment, sublicense or use shall be void;
- (ii) make error corrections to or otherwise modify or adapt the Software or create derivative works based upon the Software, or permit third parties to do the same;
- (iii) reverse engineer or decompile, decrypt, disassemble or otherwise reduce the Software to human-readable form, except to the extent otherwise expressly permitted under applicable law notwithstanding this restriction or except to the extent that Cisco is legally required to permit such specific activity pursuant to any applicable open source license;
- (iv) publish any results of benchmark tests run on the Software;
- (v) use or permit the Software to be used to perform services for third parties, whether on a service bureau or time sharing basis or otherwise, without the express written authorization of Cisco; or
- (vi) disclose, provide, or otherwise make available trade secrets contained within the Software and Documentation in any form to any third party without the prior written consent of Cisco. Customer shall implement reasonable security measures to protect such trade secrets.

To the extent required by applicable law, and at Customer's written request, Cisco shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of Cisco's applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Cisco makes such information available.

**Software, Upgrades and Additional Copies.** NOTWITHSTANDING ANY OTHER PROVISION OF THE AGREEMENT: (1) CUSTOMER HAS NO LICENSE OR RIGHT TO MAKE OR USE ANY ADDITIONAL COPIES OR UPGRADES UNLESS CUSTOMER, AT THE TIME OF MAKING OR ACQUIRING SUCH COPY OR UPGRADE, ALREADY HOLDS A VALID LICENSE TO THE ORIGINAL SOFTWARE AND HAS PAID THE APPLICABLE FEE TO AN APPROVED SOURCE FOR THE UPGRADE OR ADDITIONAL COPIES; (2) USE OF UPGRADES IS LIMITED TO CISCO EQUIPMENT SUPPLIED BY AN APPROVED SOURCE FOR WHICH CUSTOMER IS THE ORIGINAL END USER PURCHASER OR LESSEE OR OTHERWISE HOLDS A VALID LICENSE TO USE THE SOFTWARE WHICH IS BEING UPGRADED; AND (3) THE MAKING AND USE OF ADDITIONAL COPIES IS LIMITED TO NECESSARY BACKUP PURPOSES ONLY.

**Proprietary Notices.** Customer agrees to maintain and reproduce all copyright, proprietary, and other notices on all copies, in any form, of the Software in the same form and manner that such copyright and other proprietary notices are included on the Software. Except as expressly authorized in the Agreement, Customer shall not make any copies or duplicates of any Software without the prior written permission of Cisco.

**Term and Termination.** The Agreement and the license granted herein shall remain effective until terminated. Customer may terminate the Agreement and the license at any time by destroying all copies of Software and any Documentation. Customer's rights under the Agreement will terminate immediately without notice from Cisco if Customer fails to comply with any provision of the Agreement. Upon termination, Customer shall destroy all copies of Software and Documentation in its possession or control. All confidentiality obligations of Customer, all restrictions and limitations imposed on the Customer under the section titled "General Limitations" and all limitations of liability and disclaimers and restrictions of warranty shall survive termination of this Agreement. In addition, the provisions of the sections titled "U.S. Government End User Purchasers" and "General Terms Applicable to the Limited Warranty Statement and End User License Agreement" shall survive termination of the Agreement.

**Customer Records.** Customer grants to Cisco and its independent accountants the right to examine Customer's books, records and accounts during Customer's normal business hours to verify compliance with this Agreement. In the event such audit discloses non-compliance with this Agreement, Customer shall promptly pay to Cisco the appropriate license fees, plus the reasonable cost of conducting the audit.

**Export, Re-Export, Transfer and Use Controls.** The Software, Documentation and technology or direct products thereof (hereafter referred to as Software and Technology), supplied by Cisco under the Agreement are subject to export controls under the laws and regulations of the United States (U.S.) and any other applicable countries' laws and regulations. Customer shall comply with such laws and regulations governing export, re-export, transfer and use of Cisco Software and Technology and will obtain all required U.S. and local authorizations, permits, or licenses. Cisco and Customer each agree to provide the other information, support documents, and assistance as may reasonably be required by the other in connection with securing authorizations or licenses. Information regarding compliance with export, re-export, transfer and use may be located at the following URL:

[http://www.cisco.com/web/about/doing\\_business/legal/global\\_export\\_trade/general\\_export/contract\\_compliance.html](http://www.cisco.com/web/about/doing_business/legal/global_export_trade/general_export/contract_compliance.html).

**U.S. Government End User Purchasers.** The Software and Documentation qualify as "commercial items," as that term is defined at Federal Acquisition Regulation ("FAR") (48 C.F.R.) 2.101, consisting of "commercial computer software" and "commercial computer software documentation" as such terms are used in FAR 12.212. Consistent with FAR 12.212 and DoD FAR Supp. 227.7202-1 through 227.7202-4, and notwithstanding any other FAR or other contractual clause to the contrary in any agreement into which the Agreement may be incorporated, Customer may provide to Government end user or, if the Agreement is direct, Government end user will acquire, the Software and Documentation with only those rights set forth in the Agreement. Use of either the Software or Documentation or both constitutes agreement by the Government that the Software and Documentation are "commercial computer software" and "commercial computer software documentation," and constitutes acceptance of the rights and restrictions herein.

**Identified Components; Additional Terms.** The Software may contain or be delivered with one or more components, which may include third-party components, identified by Cisco in the Documentation, readme.txt file, third-party click-accept or elsewhere (e.g. on [www.cisco.com](http://www.cisco.com)) (the "Identified Component(s)") as being subject to different license agreement terms, disclaimers of warranties, limited warranties or other terms and conditions (collectively, "Additional Terms") than those set forth herein. You agree to the applicable Additional Terms for any such Identified Component(s)."

### Limited Warranty

Subject to the limitations and conditions set forth herein, Cisco warrants that commencing from the date of shipment to Customer (but in case of resale by an Approved Source other than Cisco, commencing not more than ninety (90) days after original shipment by Cisco), and continuing for a period of the longer of (a) ninety (90) days or (b) the warranty period (if any) expressly set forth as applicable specifically to software in the warranty card accompanying the product of which the Software is a part (the "Product") (if any): (a) the media on which the Software is furnished will be free of defects in materials and workmanship under normal use; and (b) the Software substantially conforms to the Documentation. The date of shipment of a Product by Cisco is set forth on the packaging material in which the Product is shipped. Except for the foregoing, the Software is provided "AS IS". This limited warranty extends only to the Software purchased from an Approved Source by a Customer who is the first registered end user. Customer's sole and exclusive remedy and the entire liability of Cisco and its suppliers under this limited warranty will be (i) replacement of defective media and/or (ii) at Cisco's option, repair, replacement, or refund of the purchase price of the Software, in both cases subject to the condition that any error or defect constituting a breach of this limited warranty is reported to the Approved Source supplying the Software to Customer, within the warranty period. Cisco or the Approved Source supplying the Software to Customer may, at its option, require return of the Software and/or Documentation as a condition to the remedy. In no event does Cisco warrant that the Software is

error free or that Customer will be able to operate the Software without problems or interruptions. In addition, due to the continual development of new techniques for intruding upon and attacking networks, Cisco does not warrant that the Software or any equipment, system or network on which the Software is used will be free of vulnerability to intrusion or attack.

**Restrictions.** This warranty does not apply if the Software, Product or any other equipment upon which the Software is authorized to be used (a) has been altered, except by Cisco or its authorized representative, (b) has not been installed, operated, repaired, or maintained in accordance with instructions supplied by Cisco, (c) has been subjected to abnormal physical or electrical stress, abnormal environmental conditions, misuse, negligence, or accident; or (d) is licensed for beta, evaluation, testing or demonstration purposes. The Software warranty also does not apply to (e) any temporary Software modules; (f) any Software not posted on Cisco's Software Center; (g) any Software that Cisco expressly provides on an "AS IS" basis on Cisco's Software Center; (h) any Software for which an Approved Source does not receive a license fee; and (i) Software supplied by any third party which is not an Approved Source.

#### **DISCLAIMER OF WARRANTY**

**EXCEPT AS SPECIFIED IN THIS WARRANTY SECTION, ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS, AND WARRANTIES INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTY OR CONDITION OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, SATISFACTORY QUALITY, NON-INTERFERENCE, ACCURACY OF INFORMATIONAL CONTENT, OR ARISING FROM A COURSE OF DEALING, LAW, USAGE, OR TRADE PRACTICE, ARE HEREBY EXCLUDED TO THE EXTENT ALLOWED BY APPLICABLE LAW AND ARE EXPRESSLY DISCLAIMED BY CISCO, ITS SUPPLIERS AND LICENSORS. TO THE EXTENT THAT ANY OF THE SAME CANNOT BE EXCLUDED, SUCH IMPLIED CONDITION, REPRESENTATION AND/OR WARRANTY IS LIMITED IN DURATION TO THE EXPRESS WARRANTY PERIOD REFERRED TO IN THE "LIMITED WARRANTY" SECTION ABOVE. BECAUSE SOME STATES OR JURISDICTIONS DO NOT ALLOW LIMITATIONS ON HOW LONG AN IMPLIED WARRANTY LASTS, THE ABOVE LIMITATION MAY NOT APPLY IN SUCH STATES. THIS WARRANTY GIVES CUSTOMER SPECIFIC LEGAL RIGHTS, AND CUSTOMER MAY ALSO HAVE OTHER RIGHTS WHICH VARY FROM JURISDICTION TO JURISDICTION.** This disclaimer and exclusion shall apply even if the express warranty set forth above fails of its essential purpose.

**Disclaimer of Liabilities - Limitation of Liability.** IF YOU ACQUIRED THE SOFTWARE IN THE UNITED STATES, LATIN AMERICA, CANADA, JAPAN OR THE CARIBBEAN, NOTWITHSTANDING ANYTHING ELSE IN THE AGREEMENT TO THE CONTRARY, ALL LIABILITY OF CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS COLLECTIVELY, TO CUSTOMER, WHETHER IN CONTRACT, TORT (INCLUDING NEGLIGENCE), BREACH OF WARRANTY OR OTHERWISE, SHALL NOT EXCEED THE PRICE PAID BY CUSTOMER TO ANY APPROVED SOURCE FOR THE SOFTWARE THAT GAVE RISE TO THE CLAIM OR IF THE SOFTWARE IS PART OF ANOTHER PRODUCT, THE PRICE PAID FOR SUCH OTHER PRODUCT. THIS LIMITATION OF LIABILITY FOR SOFTWARE IS CUMULATIVE AND NOT PER INCIDENT (I.E. THE EXISTENCE OF TWO OR MORE CLAIMS WILL NOT ENLARGE THIS LIMIT).

IF YOU ACQUIRED THE SOFTWARE IN EUROPE, THE MIDDLE EAST, AFRICA, ASIA OR OCEANIA, NOTWITHSTANDING ANYTHING ELSE IN THE AGREEMENT TO THE CONTRARY, ALL LIABILITY OF CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS COLLECTIVELY, TO CUSTOMER, WHETHER IN CONTRACT, TORT (INCLUDING NEGLIGENCE), BREACH OF WARRANTY OR OTHERWISE, SHALL NOT EXCEED THE PRICE PAID BY CUSTOMER TO CISCO FOR THE SOFTWARE THAT GAVE RISE TO THE CLAIM OR IF THE SOFTWARE IS PART OF ANOTHER PRODUCT, THE PRICE PAID

FOR SUCH OTHER PRODUCT. THIS LIMITATION OF LIABILITY FOR SOFTWARE IS CUMULATIVE AND NOT PER INCIDENT (I.E. THE EXISTENCE OF TWO OR MORE CLAIMS WILL NOT ENLARGE THIS LIMIT). NOTHING IN THE AGREEMENT SHALL LIMIT (I) THE LIABILITY OF CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS TO CUSTOMER FOR PERSONAL INJURY OR DEATH CAUSED BY THEIR NEGLIGENCE, (II) CISCO'S LIABILITY FOR FRAUDULENT MISREPRESENTATION, OR (III) ANY LIABILITY OF CISCO WHICH CANNOT BE EXCLUDED UNDER APPLICABLE LAW.

***Disclaimer of Liabilities - Waiver of Consequential Damages and Other Losses.*** IF YOU ACQUIRED THE SOFTWARE IN THE UNITED STATES, LATIN AMERICA, THE CARIBBEAN OR CANADA, REGARDLESS OF WHETHER ANY REMEDY SET FORTH HEREIN FAILS OF ITS ESSENTIAL PURPOSE OR OTHERWISE, IN NO EVENT WILL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY LOST REVENUE, PROFIT, OR LOST OR DAMAGED DATA, BUSINESS INTERRUPTION, LOSS OF CAPITAL, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR PUNITIVE DAMAGES HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY OR WHETHER ARISING OUT OF THE USE OF OR INABILITY TO USE SOFTWARE OR OTHERWISE AND EVEN IF CISCO OR ITS SUPPLIERS OR LICENSORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME STATES OR JURISDICTIONS DO NOT ALLOW LIMITATION OR EXCLUSION OF CONSEQUENTIAL OR INCIDENTAL DAMAGES, THE ABOVE LIMITATION MAY NOT APPLY TO YOU.

IF YOU ACQUIRED THE SOFTWARE IN JAPAN, EXCEPT FOR LIABILITY ARISING OUT OF OR IN CONNECTION WITH DEATH OR PERSONAL INJURY, FRAUDULENT MISREPRESENTATION, AND REGARDLESS OF WHETHER ANY REMEDY SET FORTH HEREIN FAILS OF ITS ESSENTIAL PURPOSE OR OTHERWISE, IN NO EVENT WILL CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS BE LIABLE FOR ANY LOST REVENUE, PROFIT, OR LOST OR DAMAGED DATA, BUSINESS INTERRUPTION, LOSS OF CAPITAL, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR PUNITIVE DAMAGES HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY OR WHETHER ARISING OUT OF THE USE OF OR INABILITY TO USE SOFTWARE OR OTHERWISE AND EVEN IF CISCO OR ANY APPROVED SOURCE OR THEIR SUPPLIERS OR LICENSORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

IF YOU ACQUIRED THE SOFTWARE IN EUROPE, THE MIDDLE EAST, AFRICA, ASIA OR OCEANIA, IN NO EVENT WILL CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS, BE LIABLE FOR ANY LOST REVENUE, LOST PROFIT, OR LOST OR DAMAGED DATA, BUSINESS INTERRUPTION, LOSS OF CAPITAL, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR PUNITIVE DAMAGES, HOWSOEVER ARISING, INCLUDING, WITHOUT LIMITATION, IN CONTRACT, TORT (INCLUDING NEGLIGENCE) OR WHETHER ARISING OUT OF THE USE OF OR INABILITY TO USE THE SOFTWARE, EVEN IF, IN EACH CASE, CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS, HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME STATES OR JURISDICTIONS DO NOT ALLOW LIMITATION OR EXCLUSION OF CONSEQUENTIAL OR INCIDENTAL DAMAGES, THE ABOVE LIMITATION MAY NOT FULLY APPLY TO YOU. THE FOREGOING EXCLUSION SHALL NOT APPLY TO ANY LIABILITY ARISING OUT OF OR IN CONNECTION WITH: (I) DEATH OR PERSONAL INJURY, (II) FRAUDULENT MISREPRESENTATION, OR (III) CISCO'S LIABILITY IN CONNECTION WITH ANY TERMS THAT CANNOT BE EXCLUDED UNDER APPLICABLE LAW.

Customer acknowledges and agrees that Cisco has set its prices and entered into the Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the parties.

**Controlling Law, Jurisdiction.** If you acquired, by reference to the address on the purchase order accepted by the Approved Source, the Software in the United States, Latin America, or the Caribbean, the Agreement and warranties ("Warranties") are controlled by and construed under the laws of the State of California, United States of America, notwithstanding any conflicts of law provisions; and the state and federal courts of California shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. If you acquired the Software in Canada, unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of the Province of Ontario, Canada, notwithstanding any conflicts of law provisions; and the courts of the Province of Ontario shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. If you acquired the Software in Europe, the Middle East, Africa, Asia or Oceania (excluding Australia), unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of England, notwithstanding any conflicts of law provisions; and the English courts shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. In addition, if the Agreement is controlled by the laws of England, no person who is not a party to the Agreement shall be entitled to enforce or take the benefit of any of its terms under the Contracts (Rights of Third Parties) Act 1999. If you acquired the Software in Japan, unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of Japan, notwithstanding any conflicts of law provisions; and the Tokyo District Court of Japan shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. If you acquired the Software in Australia, unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of the State of New South Wales, Australia, notwithstanding any conflicts of law provisions; and the State and federal courts of New South Wales shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. If you acquired the Software in any other country, unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of the State of California, United States of America, notwithstanding any conflicts of law provisions; and the state and federal courts of California shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties.

For all countries referred to above, the parties specifically disclaim the application of the UN Convention on Contracts for the International Sale of Goods. Notwithstanding the foregoing, either party may seek interim injunctive relief in any court of appropriate jurisdiction with respect to any alleged breach of such party's intellectual property or proprietary rights. If any portion hereof is found to be void or unenforceable, the remaining provisions of the Agreement and Warranties shall remain in full force and effect. Except as expressly provided herein, the Agreement constitutes the entire agreement between the parties with respect to the license of the Software and Documentation and supersedes any conflicting or additional terms contained in any Purchase Order or elsewhere, all of which terms are excluded. The Agreement has been written in the English language, and the parties agree that the English version will govern.

Product warranty terms and other information applicable to Cisco products are available at the following URL:

<http://www.cisco.com/go/warranty>

# Supplemental End User License Agreement for Cisco Systems Content Security Software

## IMPORTANT: READ CAREFULLY

This Supplemental End User License Agreement ("SEULA") contains additional terms and conditions for the Software product licensed under the End User License Agreement ("EULA") between You ("You" as used herein means You and the business entity you represent or "Company") and Cisco (collectively, the "Agreement"). Capitalized terms used in this SEULA but not defined will have the meanings assigned to them in the EULA. To the extent that there is a conflict between the terms and conditions of the EULA and this SEULA, the terms and conditions of this SEULA will take precedence.

In addition to the limitations set forth in the EULA on your access and use of the Software, you agree to comply at all times with the terms and conditions provided in this SEULA.

DOWNLOADING, INSTALLING, OR USING THE SOFTWARE CONSTITUTES ACCEPTANCE OF THE AGREEMENT, AND YOU ARE BINDING YOURSELF AND THE BUSINESS ENTITY THAT YOU REPRESENT TO THE AGREEMENT. IF YOU DO NOT AGREE TO ALL OF THE TERMS OF THE AGREEMENT, THEN CISCO IS UNWILLING TO LICENSE THE SOFTWARE TO YOU AND (A) YOU MAY NOT DOWNLOAD, INSTALL OR USE THE SOFTWARE, AND (B) YOU MAY RETURN THE SOFTWARE (INCLUDING ANY UNOPENED CD PACKAGE AND ANY WRITTEN MATERIALS) FOR A FULL REFUND, OR, IF THE SOFTWARE AND WRITTEN MATERIALS ARE SUPPLIED AS PART OF ANOTHER PRODUCT, YOU MAY RETURN THE ENTIRE PRODUCT FOR A FULL REFUND. YOUR RIGHT TO RETURN AND REFUND EXPIRES 30 DAYS AFTER PURCHASE FROM CISCO OR AN AUTHORIZED CISCO RESELLER, AND APPLIES ONLY IF YOU ARE THE ORIGINAL END USER PURCHASER.

For purposes of this SEULA, the Product name and the Product description You have ordered is any of the following Cisco Systems Email Security Appliance ("ESA"), Cisco Systems Web Security Appliance ("WSA") and Cisco Systems Security Management Application ("SMA") (collectively, "Content Security") and their Virtual Appliance equivalent ("Software"):

Cisco AsyncOS for Email

Cisco AsyncOS for Web

Cisco AsyncOS for Management

Cisco Email Anti-Spam, Sophos Anti-Virus

Cisco Email Outbreak Filters

Cloudmark Anti-Spam

Cisco Image Analyzer

McAfee Anti-Virus

Cisco Intelligent Multi-Scan

Cisco RSA Data Loss Prevention

Cisco Email Encryption

Cisco Email Delivery Mode

Cisco Web Usage Controls

Cisco Web Reputation

Sophos Anti-Malware

Webroot Anti-Malware

McAfee Anti-Malware  
Cisco Email Reporting  
Cisco Email Message Tracking  
Cisco Email Centralized Quarantine  
Cisco Web Reporting  
Cisco Web Policy and Configuration Management  
Cisco Advanced Web Security Management with Splunk  
Email Encryption for Encryption Appliances  
Email Encryption for System Generated Bulk Email  
Email Encryption and Public Key Encryption for Encryption Appliances  
Large Attachment Handling for Encryption Appliances  
Secure Mailbox License for Encryption Appliances

### Definitions

For purposes of this SEULA, the following definitions apply:

"Company Service" means the Company's email, Internet, security management services provided to End Users for the purposes of conducting Company's internal business.

"End User" means: (1) for the WSA and SMA, the employee, contractor or other agent authorized by Company to access the Internet and the SMA via the Company Service; and (2) for the ESA, the email boxes of the employees, contractors, or other agent authorized by Company to access or use the email services via the Company Service.

"Ordering Document" means the purchase agreement, evaluation agreement, beta, pre-release agreement or similar agreement between the Company and Cisco or the Company and a Cisco reseller, or the valid terms of any purchase order accepted by Cisco in connection therewith, containing the purchase terms for the Software license granted by this Agreement.

"Personally Identifiable Information" means any information that can be used to identify an individual, including, but not limited to, an individual's name, user name, email address and any other personally identifiable information.

"Server" means a single physical computer or devices on a network that manages or provides network resources for multiple users.

"Services" means Cisco Software Subscription Services.

"Service Description" means the description of the Software Subscription Support Services at [http://www.cisco.com/web/about/doing\\_business/legal/service\\_descriptions/index.html](http://www.cisco.com/web/about/doing_business/legal/service_descriptions/index.html)

"Telemetry Data" means samples of Company's email and web traffic, including data on email message and web request attributes and information on how different types of email messages and web requests were handled by Company's Cisco hardware products. Email message metadata and web requests included in Telemetry Data are anonymized and obfuscated to remove any Personally Identifiable Information.

"Term" means the length of the Software subscription You purchased, as indicated in your Ordering Document.

"Virtual Appliance" means the virtual version of Cisco's email security appliances, web security appliances, and security management appliances.



"Virtual Machine" means a software container that can run its own operating system and execute applications like a Server.

**Additional License Terms and Conditions****LICENSE GRANTS AND CONSENT TO TERMS OF DATA COLLECTION****License of Software.**

By using the Software and the Documentation, Company agrees to be bound by the terms of this Agreement, and so long as Company is in compliance with this Agreement, Cisco hereby grants to Company a nonexclusive, non-sublicensable, non-transferable, worldwide license during the Term to use the Software only on Cisco's hardware products, or in the case of the Virtual Appliances, on a Virtual Machine, solely in connection with the provision of the Company Service to End Users. The number of End Users licensed for the use of the Software is limited to the number of End Users specified in the Ordering Documents. In the event that the number of End Users in connection with the provision of the Company Service exceeds the number of End Users specified in the Ordering Documents, Company shall contact an Approved Source to purchase additional licenses for the Software. The duration and scope of this license(s) is further defined in the Ordering Document. The Ordering Document supersedes the EULA with respect to the term of the Software license. Except for the license rights granted herein, no right, title or interest in any Software is granted to the Company by Cisco, Cisco's resellers or their respective licensors. Your entitlement to Upgrades to the Software is subject to the Service Description. This Agreement and the Services are co-terminus.

**Consent and License to Use Data.**

Subject to the Cisco Privacy Statement at <http://www.cisco.com/web/siteassets/legal/privacy.html>, Company hereby consents and grants to Cisco a license to collect and use Telemetry Data from the Company. Cisco does not collect or use Personally Identifiable Information in the Telemetry Data. Cisco may share aggregated and anonymous Telemetry Data with third parties to assist us in improving your user experience and the Software and other Cisco security products and services. Company may terminate Cisco's right to collect Telemetry Data at any time by disabling SenderBase Network Participation in the Software. Instructions to enable or disable SenderBase Network Participation are available in the Software configuration guide.

**Description of Other Rights and Obligations**

Please refer to the Cisco Systems, Inc. End User License Agreement, Privacy Statement and Service Description of Software Subscription Support Services.



---

## A

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Advanced Malware Protection</b> | File reputation and file analysis services.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Allowed Hosts</b>               | Computers that are allowed to relay email through the Cisco appliance via a private listener. Allowed hosts are defined by their hostnames or IP addresses.                                                                                                                                                                                                                                                                                                           |
| <b>Anti-Virus</b>                  | Sophos and McAfee Anti-Virus scanning engines provide cross-platform anti-virus protection, detection and disinfection. through virus detection engines which scans files for viruses, Trojan horses and worms. These programs come under the generic term of <i>malware</i> , 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

|                  |                                                                                                                                                                                         |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Blacklist</b> | 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 <code>\$BLOCKED mail flow policy</code> . |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

---

## C

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Character Set (Double-byte)</b> | Double Byte Character Sets are foreign-language character sets requiring more than one byte of information to express each character.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>CIDR Notation</b>               | 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. |
| <b>Content Filters</b>             | Content-based filters used to process messages during the Per-Recipient Scanning phase of the work queue in the email pipeline. Content filters are evoked after Message filters, and act on individual splintered messages.                                                                                                                                                                                                                                                                                                                                  |

**Content Matching Classifier**

The detection component of the RSA data loss prevention scanning engine. A classifier contains a number of rules for detecting sensitive data, along with context rules that search for supporting data. For example, a credit card classifier not only requires that the message contain a string that matches a credit card number, but that it also contains supporting information such as an expiration data, a credit card company name, or an address.

**Conversational Bounce**

A bounce that occurs within the SMTP conversation. The two types of conversational bounces are *hard bounces* and *soft bounces*.

---

**D**
**Debounce Timeout**

The amount of time, in seconds, the system will refrain from sending the identical alert to the user.

**Delayed Bounce**

A bounce that occurs within the SMTP conversation. The recipient host accepts the message for delivery, only to bounce it at a later time.

**Delivery**

The act of delivering email messages to recipient domains or internal mail hosts from the Cisco appliance from a specific IP interface. The Cisco appliance can deliver messages from multiple IP interfaces within same physical machine using Virtual Gateway technology. Each Virtual Gateway contains a distinct IP address, hostname and domain, and email queue, and you can configure different mail flow policies and scanning strategies for each.

You can tailor the configuration of the delivery that the Cisco appliance performs, including the maximum simultaneous connections to remote hosts, the per-Virtual Gateway limit of maximum simultaneous connections to the host, and whether the conversations to remote hosts are encrypted.

**DLP**

Data loss prevention. RSA Security DLP scanning engine protects your organization's information and intellectual property and enforces regulatory and organizational compliance by preventing users from unintentionally emailing sensitive data.

**DLP Incident**

A data loss prevention incident occurs when a DLP policy detects one or more DLP violations that merit attention in an outgoing message.

**DLP Policy**

A data loss prevention policy is a set of conditions used to determine whether an outgoing message contains sensitive data and the actions that AsyncOS takes on a message that contains such data.

**DLP Risk Factor**

A score of 0 to 100 that represents the security risk of the DLP violations detected in an outgoing message. Based on the risk factor, the DLP policy determines the actions to take on the message.

**DLP Violation**

An instance of data being found in a message that violates your organization's DLP rules.

**DNS**

Domain Name System. See RFC 1045 and RFC 1035. DNS servers on a network resolve IP addresses to hostnames, and vice versa.

**DoS attack** Denial of Service attack, can also be in the form of DDos (Distributed Denial of Service Attack). An attack on a network or computer, the primary aim of which is to disrupt access to a given service.

**DSN** Delivery Status Notification, a bounced message.

---

## E

**Email Security Manager** A single, comprehensive dashboard to manage all email security services and applications on IronPort appliances. Email Security Manager allows you to manage Outbreak Filters, Anti-Spam, Anti-Virus, and email content policies — on a per-recipient or per-sender basis, through distinct inbound and outbound policies. See also *Content Filters*.

**Envelope Recipient** The recipient of an email message, as defined in the RCPT TO: SMTP command. Also sometimes referred to as the “Recipient To” or “Envelope To” address.

**Envelope Sender** The sender of an email message, as defined in the MAIL FROM: SMTP command. Also sometimes referred to as the “Mail From” or “Envelope From” address.

---

## F

**False Negative** A spam message or a message containing a virus or a DLP violation that was not detected as such.

**False Positive** A message falsely categorized as spam or as containing a virus or DLP violation.

**Fully-Qualified Domain Name (FQDN)** A domain name including all higher level domain names up to the top-level domain name; for example: mail3.example.com is a fully qualified domain name for the *host* at 192.168.42.42; example.com is the fully qualified domain name for the example.com *domain*. The fully qualified domain name must be unique within the Internet.

---

## H

**Hard Bounced Message** A message that is permanently undeliverable. This can happen during the SMTP conversation or afterward.

**HAT** Host Access Table. The HAT maintains a set of rules that control incoming connections from remote hosts for a listener. Every *listener* has its own HAT. HATs are defined for public and private listeners, and contain *mail flow policies* and *sender groups*.

---

**I**

**IDE File** Virus Definition File. An IDE file contains signatures or definitions used by anti-virus software to detect viruses.

---

**L**

**LDAP** Lightweight Directory Access Protocol. A protocol used to access information about people (including email addresses), organizations, and other resources in an Internet directory or intranet directory.

**Listener** A listener describes an email processing service that will be configured on a particular IP interface. Listeners only apply to email entering the Cisco appliance — either from the internal systems within your network or from the Internet. IronPort AsyncOS uses listeners to specify criteria that messages must meet in order to be accepted and relayed to recipient hosts. You can think of a listener as an “email injector” or even a “SMTP daemon” running for each IP address you specify.

IronPort AsyncOS differentiates between *public* listeners — which by default have the characteristics for receiving email from the Internet — and *private* listeners that are intended to accept email only from internal (groupware, POP/IMAP, and other message generation) systems.

**Log Subscription** Creation of log files that monitor the performance of the Cisco appliance. The log files are stored in local disk(s) and can also be transferred and stored in a remote system. Typical attributes of a log subscription include: name, component to monitor (email operations, server), format, and transfer method.

---

**M**

**Mail Flow Policies** A mail flow policy is a way of expressing a group of *Host Access Table* (HAT) parameters (an access rule, followed by *rate limiting* parameters and custom SMTP codes and responses) for a *listener*. Together, *sender groups* and mail flow policies are defined in a listener’s HAT. Your Cisco appliance ships with the predefined mail flow policies and sender groups for listeners.

**MAIL FROM** See *Envelope Sender*.

**Maximum Number of Retries** The maximum number of times that redelivery of a *soft bounced* message will be attempted before being *hard bounced*.

**Maximum Time in Queue** The maximum length of time that a *soft bounced* message will stay in the email queue for *delivery* before being *hard bounced*.

|                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>MTA</b>       | Mail Transfer Agent, or Messaging Transfer Agent. The program responsible for accepting, routing, and delivering email messages. Upon receiving a message from a Mail User Agent or another MTA, the MTA stores a message temporarily locally, analyses the recipients, and routes it to another MTA (routing). It may edit and/or add to the message headers. The Cisco 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. |
| <b>MUA</b>       | 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.                                                                                                                                                                                                                                                                                                                                              |
| <b>MX Record</b> | Specifies the MTA on the Internet responsible for accepting mail for a specified domain. A Mail Exchange record creates a mail route for a domain name. A domain name can have multiple mail routes, each assigned a priority number. The mail route with the lowest number identifies the primary server responsible for the domain. Other mail servers listed will be used as backup.                                                                                                                                                                                            |

---

## N

|                                  |                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Non-Conversational Bounce</b> | 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). |
| <b>NTP</b>                       | Network Time Protocol. The <code>ntpconfig</code> command configures IronPort AsyncOS to use Network Time Protocol (NTP) to synchronize the system clock with other computers.                                                                                                                                                                                     |

---

## O

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Open Relay</b>       | 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 <i>spam</i> ) through your <i>gateway</i> . The <code>listenerconfig</code> and <code>systemsetup</code> commands prevent you from unintentionally configuring your system as an open relay. |
| <b>Outbreak Filters</b> | IronPort’s Outbreak Filters feature provides an additional layer of protection from viruses. The Outbreak Filters feature quarantines suspicious email messages, holding the messages until an updated virus IDE is available. or until they are deemed not a threat.                                                                                                                                                                                                                                                          |

---

**Q****Queue**

In the Cisco appliance, you can delete, bounce, suspend, or redirect messages in the email queue. This email queue of messages for destination domains is also referred to as the *delivery queue*. The queue of messages waiting to be processed by IronPort Anti-Spam or message filter actions is referred to as the *work queue*. You can view the status of both queues using the `status detail` command.

---

**R****RAT**

Recipient Access Table. The Recipient Access Table defines which recipients will be accepted by a public listener. The table specifies the address (which may be a partial address or hostname) and whether to accept or reject it. You can optionally include the SMTP response to the RCPT TO command for that recipient. The RAT typically contains your local domains.

**Rate Limiting**

Rate limiting limits the maximum number of messages per session, the maximum number of recipients per message, the maximum message size, the maximum recipients per hour, and the maximum number of concurrent connections you are willing to accept from a remote host.

**RCPT TO**

See *Envelope Recipient*.

**Receiving**

The act of receiving email messages on a specific listener configured on an IP interface. The Cisco appliance configures listeners to receive email messages — either inbound from the Internet, or outbound from your internal systems.

**Reputation Filter**

A way of filtering suspicious senders based on their reputation. The SenderBase Reputation Service provides an accurate, flexible way for you to reject or “throttle” suspected *spam* based on the connecting IP address of the remote host.

---

**S****Sender Group**

A sender group is simply a list of senders gathered together for the purposes of handling email from those senders in the same way (that is, applying a mail flow policy to a group of senders). A sender group is a list of senders (identified by IP address, IP range, host/domain, SenderBase Reputation Service classification, SenderBase Reputation score range, or DNS List query response) separated by commas in a listener’s Host Access Table (HAT). You assign a name for sender groups, as well as *mail flow policies*.

**Soft Bounced Message**

A message whose delivery will be reattempted at a later time base on the configured *maximum number of retries* or *maximum time in queue*.

**Spam** Unwanted, Unsolicited Commercial bulk Email (UCE/UBE). Anti-spam scanning identifies email messages that are suspected to be spam, according to its filtering rules.

**STARTTLS** Transport Layer Security (TLS) is an improved version of the Secure Socket Layer (SSL) technology. It is a widely used mechanism for encrypting SMTP conversations over the Internet. The IronPort AsyncOS operating system supports the STARTTLS extension to SMTP (Secure SMTP over TLS), described in RFC 2487.

---

## T

**TOC** Threat Operations Center. This refers to all the staff, tools, data and facilities involved in detecting and responding to virus outbreaks.

---

## W

**Whitelist** A list of known good senders. Add senders you trust to the Whitelist *sender group*. The \$TRUSTED mail flow policy is configured so that email from senders you trust has no *rate limiting* enabled, and the content from those senders is not subject to anti-spam scanning.







---

## Symbols

) [13-23](#)  
/dev/null, in alias tables [24-3, 24-8](#)  
/etc/mail/aliases [24-7](#)  
/etc/mail/genericstable [24-17](#)  
\$ACCEPTED mail flow policy [7-12](#)  
\$BLOCKED mail flow policy [7-11, 7-12](#)  
\$EnvelopeSender variable [7-30](#)  
\$RELAYED mail flow policy [7-12](#)  
\$THROTTLED mail flow policy [7-11](#)  
\$TRUSTED mail flow policy [7-11, 12-13](#)

---

## Numerics

4XX error code [24-36](#)  
5XX error code [24-36](#)  
5XX SMTP response [7-11](#)

---

## A

accepting email [7-2](#)  
access privileges for custom user roles [32-9](#)  
access rules  
    in HAT [7-8](#)  
Account Privileges page [32-8](#)  
Active Directory [25-21](#)  
Active Directory Wizard [3-22](#)  
active sessions [32-30](#)  
Adaptive Scanning [14-13](#)  
address lists [7-22](#)  
    creating [7-22](#)  
    sender rate limit exceptions [7-17](#)

address literals [5-10](#)  
address rewriting [24-7](#)  
address tagging key  
    purging [24-55](#)  
admin password  
    changing [3-16, 3-25](#)  
Advanced Malware Protection [16-1](#)  
alertlisting [33-38](#)  
alert messages [3-16, 3-35](#)  
alerts  
    enabling for Outbreak Filters [14-14](#)  
    severities [33-34](#)  
alert settings [3-16, 3-35](#)  
aliasconfig command [24-8, 24-11](#)  
Alias tables  
    virtusertable [24-7](#)  
alias tables  
    aliasconfig command [24-8](#)  
    commenting [24-8](#)  
    configuring via CLI [24-7](#)  
    definition [24-7](#)  
    multiple entries [24-8](#)  
ALL entry  
    in HAT [7-2, 7-4](#)  
    in masquerading [24-18](#)  
    in RAT [8-2](#)  
alternate address [12-1](#)  
alternate MX host [24-2](#)  
altsrchost command [24-17, 24-61](#)  
always rule [14-8](#)  
AMP. See Advanced Malware Protection.  
AMP Archive [38-3](#)  
AMP Engine Logs [38-3](#)

## anti-spam

- HAT entry [7-18](#)
- HAT parameter [5-14](#)
- IronPort Anti-Spam [13-3](#)
- positive spam threshold [13-8](#)
- reporting false positives and negatives [13-15](#)
- scanning appliance-generated messages [13-14](#)
- scanning for large messages [13-5](#)
- selecting a default scanning engine [13-12](#)
- suspected spam threshold [13-8](#)
- testing [13-25](#)
- using multiple scanning engines [12-2](#)
- X-IPASFiltered header [13-5](#)

Anti-Spam Archive Logs [38-3](#)Anti-spam logs [38-3](#)anti-virus [21-20](#)

- actions [12-10](#)
- add custom header [12-11](#)
- advanced options [12-10](#)
- archive original message [12-11](#)
- dropping attachments [12-8](#)
- Encrypted [12-9](#)
- global options [12-7](#)
- mail flow policy [7-18](#)
- modify message recipient [12-12](#)
- modify message subject [12-10](#)
- per-listener actions [12-7](#)
- scan and repair [12-8](#)
- scan only [12-8](#)
- send custom alert notification [12-12](#)
- sending default notification [12-11](#)
- send to alternate destination host [12-12](#)
- Unscannable [12-9](#)
- Virus Infected [12-9](#)

Anti-Virus Archive Logs [38-3](#)Anti-Virus Logs [38-3](#)

anti-virus quarantine. See quarantine, virus

antivirus subcommand [12-7](#)appending domains [5-10](#)archivemessage command [34-34](#)archiving reports [28-35](#)AsyncOS reversion [33-30](#)AsyncOS update servers [33-22](#)AsyncOS upgrades [33-25](#)auto delivery feature [24-56](#)

## automatic update

- interval [33-23](#)

automatic updates [33-23](#)auto-select [24-56](#)AutoSupport feature [3-17, 3-35, 33-34](#)available upgrades [33-27](#)


---

**B**
bare addresses [5-10](#)Base DN [25-13](#)base entropy value, for password strength [32-20](#)blackhole listener [5-3, 40-12](#)BLACKLIST sender group [7-11](#)body scanning [9-31](#)bounceconfig Command [24-40](#)Bounce Logs [38-2](#)bounce profile [24-40](#)bouncerecipients command [34-24](#)

## bounces

- conversational [24-36](#)
- non-conversational [24-36](#)

Bounce Verification [24-51](#)

## bouncing recipients

- all [34-26](#)
- by Envelope From [34-25](#)
- by hostname [34-25](#)

## bypassing

- anti-spam [9-72](#)
- throttling [8-5](#)

## C

- call-ahead SMTP server [22-1](#)
  - routing [22-7](#)
- canonicalization [20-5](#)
- case-sensitivity
  - in CLI [2-5](#)
  - in LDAP queries [25-13, 25-18](#)
  - in message filters [9-19](#)
  - systemsetup command [3-26](#)
- categories
  - adult [15-14](#)
  - advertisements [15-14](#)
  - alcohol [15-14](#)
  - arts [15-14](#)
  - astrology [15-14](#)
  - auctions [15-14](#)
  - business and industry [15-15](#)
  - chat and instant messaging [15-15](#)
  - cheating and plagiarism [15-15](#)
  - child abuse content [15-15](#)
  - computers and internet [15-15](#)
  - computer security [15-15](#)
  - dating [15-15](#)
  - digital postcards [15-15](#)
  - dining and drinking [15-15](#)
  - dynamic and residential [15-15](#)
  - education [15-16](#)
  - entertainment [15-16](#)
  - extreme [15-16](#)
  - fashion [15-16](#)
  - file transfer services [15-16](#)
  - filter avoidance [15-16](#)
  - finance [15-16](#)
  - freeware and shareware [15-16](#)
  - gambling [15-16](#)
  - games [15-16](#)
  - government and law [15-17](#)
  - hacking [15-17](#)
  - hate speech [15-17](#)
  - health and nutrition [15-17](#)
  - humor [15-17](#)
  - illegal activities [15-17](#)
  - illegal downloads [15-17](#)
  - illegal drugs [15-18](#)
  - infrastructure and content delivery networks [15-18](#)
  - internet telephony [15-18](#)
  - job search [15-18](#)
  - lingerie and swimsuits [15-18](#)
  - lotteries [15-18](#)
  - mobile phones [15-18](#)
  - nature [15-18](#)
  - news [15-18](#)
  - non-governmental organizations [15-18](#)
  - non-sexual nudity [15-18](#)
  - online communities [15-19](#)
  - online storage and backup [15-19](#)
  - online trading [15-19](#)
  - organizational email [15-19](#)
  - parked domains [15-19](#)
  - peer file transfer [15-19](#)
  - personal sites [15-19](#)
  - photo searches and images [15-19](#)
  - politics [15-19](#)
  - pornography [15-19](#)
  - professional networking [15-19](#)
  - real estate [15-19](#)
  - reference [15-20](#)
  - religion [15-20](#)
  - SaaS and B2B [15-20](#)
  - safe for kids [15-20](#)
  - science and technology [15-20](#)
  - search engines and portals [15-20](#)
  - sex education [15-20](#)
  - shopping [15-20](#)
  - social networking [15-20](#)
  - social science [15-20](#)
  - society and culture [15-20](#)

- software updates [15-20](#)
- sports and recreation [15-20](#)
- streaming audio [15-20](#)
- streaming video [15-20](#)
- tobacco [15-21](#)
- transportation [15-21](#)
- travel [15-21](#)
- unclassified [15-21](#)
- weapons [15-21](#)
- web-based email [15-21](#)
- web hosting [15-21](#)
- web page translation [15-21](#)
- Centralized Management
  - and centralized quarantines [42-6, 42-8](#)
  - and Destination Controls [39-29](#)
  - and quarantines [30-10](#)
- centralized management [33-17](#)
- certificate
  - certificate authority [23-2](#)
- certificates
  - adding [23-3](#)
  - certificate authorities list [23-16](#)
  - demo [3-26](#)
  - DLP with RSA Enterprise Manager [17-26](#)
  - exporting [23-6](#)
  - FIPS management [27-4](#)
  - generating and signing your own [23-2](#)
  - generating a request [19-8, 23-5](#)
  - importing [23-1](#)
  - intermediate certificates [23-3](#)
- Certificate Signing Request [23-2](#)
- chain query
  - creating [25-28](#)
  - LDAP [25-28](#)
- chains, of aliases [24-8](#)
- Change Password link [32-16](#)
- changing your password [32-16](#)
- charset [31-6](#)
- CIDR address block [7-4](#)
- Cisco Content Security Management Appliance. See Security Management appliance
- Cisco Security Intelligence Operations [14-3](#)
- Cisco Web Security Services [15-3](#)
- clean message [28-8](#)
- clear command [2-8](#)
- CLI
  - configuring languages [33-61](#)
  - see *Command Line Interface*
- CLI Audit Logs [38-2](#)
- closing unsuccessful or unproductive inbound connections [5-6](#)
- cluster [15-4](#)
- command completion [2-6](#)
- command line interface (CLI)
  - case-sensitivity in [2-5](#)
  - command completion in [2-6](#)
  - conventions [2-3](#)
  - default setting [2-4](#)
  - exit [2-6](#)
  - history [2-6](#)
  - subcommands [2-5](#)
  - white space [2-5](#)
- comments [7-22, 24-5](#)
  - comments in imported files [7-22, 24-5](#)
- commit command [2-7](#)
- community string [34-37](#)
- configuration, testing [3-36](#)
- configuration file [33-7](#)
  - CLI [33-12](#)
  - GUI [33-8](#)
  - XML [33-8](#)
- conformance level
  - SPF/SIDF verification [20-24](#)
- connectivity issues, troubleshooting [40-16](#)
- content dictionary [21-1](#)
- content filters [30-2](#)
  - actions [11-9](#)
  - applied during email pipeline [11-1](#)

- conditions [11-2](#)
- example [C-12, C-13, C-14](#)
- non-ascii character sets [11-19, C-18](#)
- variables [11-14](#)
- content matching classifier [17-10, 17-15](#)
- conversational bounces [24-36](#)
- counters [34-2](#)
- CPU usage [34-4](#)
- CRAM-MD5 [25-36](#)
- CSV data [28-33](#)
- custom DLP dictionaries [17-17](#)
- custom header [13-19](#)
- custom SMTP response
  - variable [7-30](#)
- custom user roles [32-7](#)

## D

- daily magnitude [28-13](#)
- data loss prevention [30-2](#)
  - see *DLP*
- default
  - domain [8-1](#)
  - gateway [3-17, 3-26](#)
  - hostname [3-16, 3-25](#)
  - IP address [3-13](#)
  - router [3-17, 3-26](#)
  - sender domain [5-10](#)
- default DNS server [33-55](#)
- default router [3-17](#)
- defining
  - user preferences [33-61](#)
- delayed bounces [24-36](#)
- delegated administration [32-7](#)
- delete all messages in the spam quarantine [31-24](#)
- deleterecipients command [34-22](#)
- delivering mail [24-42](#)
  - controlling [24-42](#)
  - controlling mail to destination domain [24-42](#)
  - message time out [24-56](#)
  - possible delivery [24-56](#)
- delivernow command [34-31](#)
- delivery
  - encrypting [23-2](#)
- deliveryconfig command [24-57](#)
- Delivery Connection ID (DCID) [34-4](#)
- Delivery Logs [38-2](#)
- delivery mode [41-1](#)
- delivery queue [34-22](#)
- delivery queue, monitoring [34-15](#)
- Delivery Status Details page [28-16](#)
- Delivery Status page [28-15](#)
- demo certificate [3-26, 23-3](#)
- demonstration certificate [23-9](#)
- destconfig command [23-12, 24-45](#)
- Destination Controls [24-45](#)
  - and Centralized Management [39-29](#)
  - importing and exporting configurations [24-47](#)
- detection rule [17-14, 17-15, 17-19](#)
- DHAP
  - mail flow policy [7-17](#)
- diagnostic -> network -> arpshow command [40-18](#)
- diagnostic -> network -> flush command [40-18](#)
- diagnostic -> network -> smtppping command [40-19](#)
- Directory Harvest Attack (DHA) [25-29](#)
- Direct Server Return (DSR) [37-13](#)
- disclaimers
  - adding to messages [21-13](#)
  - HTML text resources [21-11](#)
  - using text resources [21-12](#)
- disclaimer stamping [21-13](#)
  - multiple encodings [21-16](#)
- DKIM
  - DNS TXT record [20-4](#)
  - domain profile [20-2](#)
  - enabled on a mail flow policy [20-2](#)
  - signing [20-2](#)
- DKIM verification [20-21](#)

- Authentication-Results header [20-21](#)
- DLP [17-1](#)
  - Assessment Wizard [17-7](#)
  - content matching classifier [17-15](#)
  - dictionaries [17-17](#)
  - exporting DLP policies [17-31](#)
  - false positives, minimizing [17-2, 17-10, 17-12, 17-13, 17-15, 17-19](#)
  - fingerprinting [17-26](#)
  - including sensitive content in Message Tracking [17-38](#)
  - message actions [17-34](#)
  - policy customizations [17-34](#)
  - quarantines [17-42](#)
  - regular expressions [17-15](#)
  - reporting [17-42](#)
  - risk factor score [17-18](#)
  - RSA Email DLP [17-4](#)
  - RSA Enterprise Manager [17-23](#)
  - severity scale [17-21](#)
  - switching modes [17-33](#)
  - troubleshooting [17-42](#)
  - updating the engine and classifiers [17-39](#)
- DLP policies
  - arranging the order [17-21](#)
  - content matching classifier [17-10](#)
  - custom policies [17-9](#)
  - detection rule [17-14, 17-15, 17-19](#)
  - filtering attachments [17-20](#)
  - filtering senders and recipients [17-20](#)
  - regular expressions [17-15](#)
  - RSA Email DLP [17-6](#)
  - severity scale [17-21](#)
  - templates [17-6, 17-8](#)
- D-Mode [17-4, 41-1](#)
- DNS [D-1](#)
  - A Record [40-22](#)
  - authoritative server [33-54](#)
  - cache [40-21](#)
  - disabling reverse DNS lookup timeout [33-56](#)
  - double lookup [7-3, 7-28, 28-11](#)
  - priority [33-55](#)
  - PTR record [40-22](#)
  - servers [3-17, 3-26](#)
  - setting [3-17, 3-26](#)
  - splitting [33-55](#)
  - testing [40-20](#)
  - timeout [33-55](#)
  - timeout for reverse DNS lookups [33-56](#)
- DNSBL [9-34](#)
- DNS cache [34-20](#)
- DNS cache, flushing [33-56](#)
- dnsconfig command [33-54](#)
- dnsflush command [33-56](#)
- DNS list [9-34](#)
- DNS lookup [34-20](#)
- DNS servers [33-54](#)
- DNS settings [33-56](#)
- dnsstatus command [34-20](#)
- DNS TXT record [20-2](#)
- domain
  - adding a default domain [5-10](#)
- domain context
  - in alias tables [24-8, 24-11](#)
- Domain Debug Logs [38-2](#)
- Domain Keys [20-1](#)
  - canonicalization [20-5](#)
  - DNS text record [20-13](#)
  - DNS TXT record [20-4](#)
  - domain profile [20-2](#)
  - enabled on a mail flow policy [20-2](#)
  - enabled via mail flow policy [7-19](#)
  - exporting domain profiles [20-14](#)
  - importing domain profiles [20-14](#)
  - importing signing keys [20-12](#)
  - selector [20-5](#)
  - signing [20-2](#)
  - signing key size [20-3](#)

- testing a domain profile [20-14](#)
    - verification [20-1](#)
    - verifying signatures [20-2](#)
  - DomainKey-Signature header [20-3](#)
  - domain map
    - commenting [24-34](#)
    - importing and exporting [24-34](#)
    - importing invalid entries [24-34](#)
    - limits [24-29](#)
    - overview [24-28](#)
  - Domain Name Service (DNS)
    - settings [3-17, 3-26](#)
  - domain profile
    - deleting all existing profiles [20-15](#)
    - exporting [20-14](#)
    - importing [20-14](#)
    - removing domain profiles [20-14](#)
    - testing [20-14](#)
  - domains [28-13](#)
  - domain table [24-28](#)
  - double-DNS verified [28-12](#)
  - drop-attachments-where-dictionary-match [9-88](#)
  - DSN (delay notification messages) [24-40](#)
  - DSR [37-13](#)
    - load balancing [37-13](#)
    - loopback interface [37-13](#)
    - Virtual IP (VIP) [37-13](#)
  - DTD (document type definition) [33-10](#)
  - dual DKIM and DomainKey signing [20-7](#)
  - dummy accounts [6-6](#)
  - duplex settings, editing [37-1](#)
  - rewriting addresses [24-7](#)
  - email address
    - source routing [5-10](#)
  - email gateway [5-1](#)
  - email injector
    - see *listener*
  - Email Security Monitor [28-1](#)
    - automated reporting [28-33](#)
    - external domains received listing [28-11](#)
    - Items Displayed menu [28-12](#)
    - summary table [28-7](#)
    - Time Range menu [28-7](#)
  - enabling DomainKeys and DKIM signing on a mail flow policy [20-2](#)
  - encoding
    - in disclaimers [21-16](#)
  - encryption [5-13, 23-1](#)
    - use with filter action [18-8, 19-13](#)
  - encryptionconfig CLI command [18-4](#)
  - encryption headers [18-11](#)
  - encryption profiles
    - configuring [18-4](#)
  - End User Quarantine
    - see spam quarantine, end user access [31-16](#)
  - End User Quarantine. See spam quarantine
  - enterprise gateway [3-37](#)
  - Enterprise Gateway configuration [5-15](#)
  - Envelope Recipient [9-25, 24-7, 29-3](#)
  - Envelope Recipient, rewriting [24-7](#)
  - Envelope Sender [9-25, 29-3](#)
  - Envelope Sender, rewriting [24-16](#)
  - envelope sender DNS verification [7-30](#)
  - Envelope To [24-7](#)
  - Envelope To, rewriting in alias tables [24-7](#)
  - Ethernet interfaces [B-1](#)
  - evaluation key
    - McAfee [3-34](#)
    - Sophos [3-34](#)
  - evaluation key for IronPort Anti-Spam [3-34, 13-3](#)
- 
- ## E
- Early Expiration
    - for quarantine [30-4](#)
  - editing DNS settings via GUI [33-56](#)
  - email
    - clean message [28-8](#)



evaluation key for McAfee [12-2](#)  
 evaluation key for Outbreak Filters [3-22, 3-35](#)  
 exception table  
     adding entries [7-35](#)  
 exit command [2-9](#)  
 explained [7-30](#)  
 exporting  
     HTML text resources [21-11](#)  
     text resources [21-10](#)  
 external authentication [25-40](#)  
     enabling LDAP [32-21](#)  
     enabling RADIUS [32-22](#)

## F

factory configuration [3-14](#)  
 feature key [33-5](#)  
 featurekey command [3-37, 12-2, 13-3](#)  
 Federal Information Processing Standard  
     see *FIPS management*  
 file analysis [16-1](#)  
 File Analysis quarantine [30-1](#)  
 file reputation filtering [16-1](#)  
 filtering unparseable messages [9-23](#)  
 filters [9-1](#)  
     comment character [9-3](#)  
     matching dictionary terms [9-15, 9-36](#)  
     matching empty headers [9-24](#)  
     regular express and Python [9-19](#)  
     scannable archive file types [9-31](#)  
     unparseable messages [9-23](#)  
 final entry, in HAT [7-2](#)  
 findevent [34-35](#)  
 finding senders [7-15](#)  
 fingerprinting [17-26](#)  
 FIPS management  
     managing certificates and keys [27-4](#)  
     overview [27-1](#)  
 firewall permissions [40-22](#)

firewall ports [D-1](#)  
 forcing updates [12-18](#)  
 forward DNS lookup [34-19](#)  
 FTP [A-1, D-1](#)  
 FTP Access [A-2](#)  
 FTP Push [38-6](#)  
 FTP Server Logs [38-2](#)  
 fully-qualified domain name [7-4](#)

## G

gateway configuration [5-1](#)  
 gauges [34-4](#)  
 genericstable file [24-19](#)  
 getting started [3-1](#)  
 global alias [24-8](#)  
 global counters [34-21](#)  
 global unsubscribe  
     adding [24-70](#)  
     commenting [24-72](#)  
     importing and exporting [24-72](#)  
     maximum entries [24-69](#)  
     overview [24-69](#)  
     syntax [24-69](#)  
 globbing [24-2](#)  
 good neighbor table [23-11](#)  
 graph [28-6](#)  
 graphical user interface  
     see *GUI*  
 graphs [36-11](#)  
 GUI  
     accessing [2-1](#)  
     browser requirements [2-1](#)  
     configuring language [33-61](#)  
     enabling [3-26, 36-7](#)  
     overview [36-7](#)  
 GUI logs. See HTTP logs  
 GUI session timeout [32-26](#)

---

**H**

hard power reset [33-29, 40-27](#)

HAT [7-15](#)

    delayed rejection [7-8](#)

    delayed rejections [5-7](#)

    exporting [7-21](#)

    importing [7-21](#)

    significant bits [7-17](#)

    testing HAT variables [7-10](#)

    using HAT variables [7-9](#)

    using HAT variables - CLI example [7-10](#)

    using HAT variables - GUI example [7-10](#)

HAT delayed rejection [7-8](#)

HAT delayed rejections [5-7](#)

HAT order

    editing via GUI [7-14](#)

headers [24-7, 24-16, 24-18](#)

    anti-spam [13-14](#)

headers, inserting [18-11](#)

headers, logging [13-24, 38-42](#)

headers, stripping with message filters [9-69](#)

help command [2-9](#)

hiding network topology [5-11, 24-16](#)

history, in CLI [2-6](#)

Host Access Table (HAT)

    definition [7-1](#)

    order in [7-2](#)

    parameters [7-8](#)

    reordering in GUI [7-14](#)

    rules [7-1](#)

    syntax [7-1](#)

Host DNS Verification, explained [7-28](#)

hostname [3-16, 3-25](#)

    specifying the hostname during setup [3-16](#)

hostname, setting [33-53](#)

hostrate command [34-17](#)

hoststatus command [28-16, 34-13](#)

HTTP [A-1, D-1](#)

    enabling [3-26](#)

    GUI [36-7](#)

HTTP authentication [28-34](#)

HTTP Logs [38-3](#)

HTTP proxy server [33-23](#)

HTTPS [A-1](#)

    certificate for [23-18](#)

    enabling [3-26](#)

    GUI [36-7](#)

HTTPS login [2-2](#)

HTTPS proxy server [33-23](#)

---

image analysis [9-80, 11-6, 11-11](#)

image scanning [9-80](#)

image verdicts [9-80](#)

IMAP authentication [31-18](#)

implementsv [7-31](#)

importing

    HTML text resources [21-11](#)

    text resources [21-10](#)

importing domain profiles [20-14](#)

importing signing keys [20-12](#)

inbound connections

    closing unsuccessful or unproductive connections [5-6](#)

inbound connection timeout [5-6](#)

Incoming Mail Reporting page [28-9](#)

incoming messages, defined [10-3](#)

Incoming Relay [13-15](#)

incoming relay

    custom header [13-19](#)

    received header [13-20](#)

Incoming Relays

    example log entry [13-24](#)

incoming relays [20-22](#)

Injection Connection ID (ICID) [34-4](#)

injection control counter reset [7-25](#)

injection control periodicity [7-25](#)  
 injection counters reset period [5-6](#)  
 Injection Debug Logs [38-2](#)  
 injector  
     see *listener*  
 insecure relay [8-2](#)  
 inserting headers [18-11](#)  
 installation [3-1](#)  
     reverting [33-30](#)  
 interface command [24-56](#)  
 invalid recipient [28-8](#)  
 IP addresses [28-14](#)  
 IP address profile pages [28-13](#)  
 IP interfaces  
     assigning [3-18, 3-25](#)  
     defining listeners on [3-27](#)  
 IPMI [34-37](#)  
 IP port  
     defining in listenerconfig command [5-8](#)  
 IronPort Anti-Spam  
     archivingY [13-9](#)  
     evaluation key [3-21, 3-34, 13-3](#)  
     filters [13-2](#)  
     testing [13-25](#)  
 IronPort Email Encryption  
     configuring [18-1](#)  
     encryption profiles [18-4](#)  
     envelope settings [18-5](#)  
     key server settings [18-5](#)  
     message settings [18-5](#)  
     notification settings [18-5](#)  
     use with filter action [18-8, 19-13](#)  
 IronPort Intelligent Multi-Scan  
     enabling [13-7](#)  
 IronPort Spam Quarantine  
     released messages and email pipeline [4-9](#)  
     stripping "SMTP:" in the LDAP query [25-43](#)  
 IronPort Spam Quarantine. See Spam quarantine  
 IronPort Text Mail Logs [38-2](#)

---

## K

keys  
     FIPS management [27-4](#)  
 key size [20-3](#)

---

## L

languages  
     defining default per user [33-61](#)  
     user preferences [33-61](#)  
 large message scanning [13-5](#)  
 last command [32-7](#)  
 LDAP [31-15, 31-16, D-2](#)  
     Accept query [5-12](#)  
     alias consolidation query [25-44](#)  
     alias expansion [25-20](#)  
     and RSA Enterprise Manager [17-30](#)  
     anonymous queries [25-14](#)  
     base DN [25-13](#)  
     chain queries [25-28](#)  
     connection pooling [25-34](#)  
     connections [25-17](#)  
     domain-based queries [25-26](#)  
     end-user authentication query [25-43](#)  
     external authentication [25-40, 32-21](#)  
     failover [25-46](#)  
     group queries [9-25, 9-26](#)  
     LDAPS certificate [25-14](#)  
     load-balancing [25-46](#)  
     mail policy [C-5](#)  
     Microsoft Exchange 5.5 support [25-9](#)  
     multiple servers [25-46](#)  
     OpenLDAP queries [25-19](#)  
     query tokens [25-13](#)  
     recursive queries [25-14](#)  
     SSL [25-14](#)  
     SunONE queries [25-20](#)  
     testing queries [25-12, 25-17](#)

- testing servers [25-6](#)
- test servers [25-6](#)
- user distinguished name query [25-45](#)
- LDAP Accept during SMTP conversation [5-12](#)
- LDAP accept query [5-12](#)
- LDAP Debug Logs [38-3](#)
- LDAP errors [25-18](#)
- LDAP routing query
  - with SMTP call-ahead recipient validation [22-6](#)
- LDAPS [D-2](#)
  - Global Catalog Server [D-2](#)
- LDAPS certificate [25-14](#)
- limits
  - altsrchoost [24-64](#)
  - SMTP Routes [24-3](#)
- link aggregation [37-3](#)
- listener
  - add a default domain [5-10](#)
  - adding disclaimers [21-13](#)
  - adding received header [5-11](#)
  - add listener [5-8](#)
  - definition [5-1](#)
  - encrypting [23-2](#)
  - encryption on [5-13](#)
  - injection counters reset period [5-6](#)
  - LDAP accept query [5-12](#)
  - loose SMTP address parsing [5-9](#)
  - malformed MAIL FROM and default domain [5-12](#)
  - maximum concurrent connections [5-6](#)
  - strict SMTP address parsing [5-9](#)
  - timeout for unsuccessful inbound connections [5-6](#)
  - Total Time Limit for All Inbound Connections [5-7](#)
- listenerconfig command [5-2](#)
- listeners
  - configuring [5-1](#)
  - private [5-1](#)
  - public [5-1](#)
- load [34-4](#)
- loadconfig command [33-14](#)
- log file type [38-1](#)
- logging
  - overview [38-1](#)
- logging,headers [13-24, 38-42](#)
- Logging Options [38-42](#)
- logheaders command [38-42](#)
- logical IP interface [3-18, 3-25](#)
- logs
  - Anti-Spam Archive [38-3](#)
  - Anti-Virus [38-3](#)
  - Anti-Virus Archive [38-3](#)
  - Bounce Logs [38-2](#)
  - CLI Audit Logs [38-2](#)
  - comparison [38-5](#)
  - Configuration History Logs [38-37](#)
  - definition [38-1](#)
  - Delivery Logs [38-2](#)
  - extensions in filenames [38-43](#)
  - format [38-1](#)
  - FTP Server Logs [38-2](#)
  - global attributes [38-40](#)
  - HTTP Logs [38-3](#)
  - Injection Debug Logs [38-2](#)
  - IronPort Text Mail Logs [38-2](#)
  - LDAP Debug Logs [38-3](#)
  - levels [38-39](#)
  - log subscription defined [38-1](#)
  - message headers in [38-42](#)
  - NTP Logs [38-3](#)
  - qmail Format Delivery Logs [38-2](#)
  - rolling over [38-7](#)
  - Scanning [38-3](#)
  - SCP Push [38-7](#)
  - Status Logs [38-2](#)
  - subscriptions [38-7](#)
  - syslog push [38-7](#)
  - troubleshooting with [40-23](#)
- log subscription [38-1](#)

IronPort Anti-Spam [13-9](#)  
 Sophos [12-11](#)  
 log subscriptions [38-7](#)  
 lookup  
   DNS A [7-3, 7-28](#)  
   DNS PTR [7-3, 7-28](#)  
 loopback interface [37-13](#)

## M

mailconfig command [3-36, 33-13](#)  
 mailertable feature [24-1](#)  
 mail flow policies  
   \$ACCEPTED [7-12](#)  
   \$BLOCKED [7-11, 7-12](#)  
   \$RELAYED [7-12](#)  
   \$THROTTLED [7-11](#)  
   \$TRUSTED [7-11](#)  
   adding via GUI [7-15](#)  
   definition of [7-8](#)  
   editing via GUI [7-13](#)  
   HAT parameters for [7-8](#)  
   predefined [7-11](#)  
 MAIL FROM [9-10, 11-8, 24-16](#)  
   configuring for notifications [33-33](#)  
 mailing lists  
   notifications [31-20](#)  
 mail loops, detecting [9-114](#)  
 mail policies  
   adding users [C-5](#)  
   First Match Wins [10-4](#)  
   LDAP [C-5](#)  
   removing users [C-6](#)  
 mail policies, outgoing  
   DLP [17-22, 17-23](#)  
   RSA Enterprise Manager [17-31](#)  
 mail protocol  
   defining in listenerconfig command [5-2](#)  
 mail transfer agent. See MTA. [42-2](#)  
 mail trend graph [28-6](#)  
 malformed entries, in alias tables [24-8](#)  
 malware  
   defined [12-2](#)  
 mapping domains [24-2](#)  
 marketing messages [28-8](#)  
 masquerade subcommand [24-19](#)  
 masquerading  
   and altsrhost command [24-17](#)  
   commenting [24-18](#)  
   configuring via CLI [24-17](#)  
   definition [24-16](#)  
   importing and exporting [24-19](#)  
   importing invalid entries [24-19](#)  
   limits [24-18](#)  
   table syntax [24-18](#)  
   via an LDAP query [24-17](#)  
   via a static table [24-17](#)  
 matched content  
   viewing [30-14](#)  
 matching empty headers [9-24](#)  
 maximum  
   concurrent connections in HAT [7-15](#)  
   message size in HAT [5-14, 7-15](#)  
   messages per connection in HAT [5-14, 7-15](#)  
   recipients per hour, in systemsetup [3-28, 3-31](#)  
   recipients per hour code in HAT [7-16](#)  
   recipients per hour exceeded text in HAT [7-16](#)  
   recipients per hour in HAT [6-7, 7-16](#)  
   recipients per message in HAT [5-14, 7-16](#)  
   recipients per time interval in HAT [7-17](#)  
 maximum concurrent connections [5-6](#)  
 maximum connections for listeners [24-57](#)  
 maximum recipients per hour [5-14](#)  
 mbox format [9-69](#)  
 mbox-format log file [12-11, 13-9](#)  
 McAfee  
   evaluation key [3-34](#)  
   update servers [33-22](#)

- McAfee anti-virus engine [12-5](#)
  - memory [34-5](#)
  - message actions
    - creating [17-34](#)
  - message actions for DLP [17-34](#)
  - message body scanning [9-31](#)
  - message encoding [5-8](#)
    - modifying [5-8, 9-98](#)
    - setting for headings and footers [5-8](#)
  - message filter
    - filter actions [9-49](#)
  - message filter action variables
    - using in disclaimers [21-15](#)
  - message filter for SBRS [6-7](#)
  - message filters [30-2](#)
    - adding [9-90](#)
    - attachment-protected [9-13](#)
    - attachment-unprotected [9-13](#)
    - body-dictionary-match [9-36](#)
    - combining [9-3, 9-16](#)
    - deleting [9-91](#)
    - encryption [9-32](#)
    - exporting [9-96](#)
    - importing [9-95](#)
    - making (in)active [9-91](#)
    - MIME types [9-31](#)
    - moving [9-91](#)
    - ordering [9-4](#)
    - overview [9-1](#)
    - random numbers in [9-30](#)
    - rules [9-2](#)
    - SenderBase Reputation Score [9-35](#)
    - status [9-92](#)
    - syntax [9-3](#)
    - time and date [9-28](#)
    - variables [9-56](#)
  - message headers [9-29, 38-42](#)
  - message headers, inserting with message filters [9-69](#)
  - Message ID (MID) [34-3](#)
  - message modification level threshold [14-19](#)
  - message replication [9-49, 9-63](#)
  - message splintering
    - defined [10-5](#)
  - Message Tracking
    - and sensitive content [17-38](#)
  - message tracking
    - Incoming Relays [13-23](#)
  - message variables
    - spam quarantine notifications [31-19](#)
  - MIB file [34-37](#)
  - Microsoft Exchange, LDAP queries for [25-21](#)
  - monitoring [34-7](#)
  - monitoring Virtual Gateway addresses [24-68](#)
  - M-Series [42-1](#)
  - M-Series appliance [42-1](#)
  - MTA [3-37, 5-1, 5-15, 23-1](#)
  - multilayer anti-virus scanning [12-2](#)
  - multiple appliances [3-14](#)
  - multiple IP interfaces [24-63](#)
  - multiple recipients [10-5](#)
  - MX [3-1](#)
  - MX records [40-22](#)
- 
- ## N
- 
- negative scores [7-6](#)
  - netmask [3-18, 3-26](#)
  - netmasks, selecting [B-1](#)
  - netstat command [40-18](#)
  - network access list [32-24](#)
  - networking worksheet [3-11](#)
  - network owner [28-13](#)
  - Network Owner profile pages [28-13](#)
  - network problems, troubleshooting [40-18](#)
  - network time protocol (NTP)
    - settings [3-16, 3-36](#)
  - network topology [B-4](#)
  - NIC pairing [37-3](#)

- alerts [37-4](#)
- named on upgrade [37-4](#)
- NIC teaming [37-3](#)
- non-conversational bounces [24-36](#)
- Normal Expiration
  - for quarantine [30-3](#)
- No Subject [29-5](#)
- not.double.verified [7-29, 7-38](#)
- nslookup command [40-20](#)
- NTP [D-2](#)
- NTP Logs [38-3](#)
- NTP server [33-59](#)
  - removing [33-60](#)
- nx.domain [7-38](#)
- NXDOMAIN [7-29, 7-38](#)

---

## O

- offline command [33-3](#)
- offline state [33-3](#)
- oldmessage command [34-34](#)
- online help [2-9](#)
- opening links in a separate window [28-7](#)
- open relay, definition [8-2](#)
- Outbreak Filters
  - Adaptive rules defined [14-7](#)
  - Adaptive Scanning [14-13](#)
  - alerts [14-23](#)
  - always rule [14-8](#)
  - anti-virus updates [14-10](#)
  - bypassed file extensions [14-18](#)
  - CASE [14-4](#)
  - Context Adaptive Scanning Engine [14-4](#)
  - delaying messages [14-4](#)
  - enabling alerts [14-14](#)
  - evaluation key [3-22, 3-35](#)
  - message modification [14-6](#)
  - modifying messages [14-6](#)
  - multiple scores [14-9](#)
  - non-viral threats [14-3](#)
  - Outbreak rules defined [14-6](#)
  - overview [14-1](#)
  - redirecting links [14-5](#)
  - re-evaluating messages [14-10, 14-11](#)
  - rule [14-7](#)
  - setting a message modification level threshold [14-19](#)
  - setting a quarantine level threshold [14-17](#)
  - skipping [11-11](#)
  - SNMP Traps [14-23](#)
  - threat categories [14-3](#)
  - updating rules [14-15](#)
  - using without anti-virus scanning [14-9](#)
  - virus outbreaks [14-3](#)
- Outgoing Destinations page [28-14](#)
- outgoing messages, defined [10-3](#)
- Outgoing Senders page [28-15](#)
- overflow [14-10](#)
- Overview page (Security Monitor) [28-6](#)

---

## P

- packet capture [40-31](#)
- partial address
  - in HAT [7-4](#)
  - in RAT [8-4](#)
- partial domain
  - in alias tables [24-8](#)
  - in Masquerading [24-18](#)
- password
  - changing [32-16](#)
  - settings [32-17](#)
- passwords, changing [32-16](#)
- pausing the work queue [34-32](#)
- PEM format, for certificates [19-8, 23-5](#)
- performance [40-26](#)
- phased approach to reputation filters [6-4](#)
- PII [17-8](#)
- ping command [40-18](#)

- pinout for serial connection [3-9](#)
- policies, predefined [7-2](#)
- POP/IMAP servers [5-15](#)
- POP authentication [31-18](#)
- positive scores [7-6](#)
- possible delivery [24-56, 24-57](#)
- preferences
  - defining for users [33-61](#)
- private injector [3-30](#)
- private key [23-1](#)
- private listeners
  - default entries [7-2](#)
- Profile for Domain pages [28-13](#)
- prototcol
  - see *mail protocol*
- proxy server [33-23](#)
- proxy server for IronPort Anti-Spam Rules [33-21](#)
- public blacklist [9-34](#)
- public listener [3-28](#)
- public listeners
  - default entries [7-2](#)
- purging address tagging keys [24-55](#)
- PVO. See quarantines, policy, virus, and outbreak

## Q

- qmail Format Delivery Logs [38-2](#)
- QMQP [D-3](#)
- quarantine [30-2](#)
  - applying actions to messages in [30-12](#)
  - default action [30-4, 30-7](#)
  - displaying non-ascii characters in subject [16-9, 30-6](#)
  - early expiration [30-4](#)
  - In other quarantines [30-13](#)
  - international character sets [30-11](#)
  - normal expiration [30-3](#)
  - outbreak [30-2](#)
  - outbreak, reporting messages to Cisco [30-18](#)
  - outbreak, special filters [30-17](#)

- retention time [30-3](#)
- spam. See Spam quarantine
- stripping attachments [16-9, 30-6](#)
- subject tagging [16-9, 30-6](#)
- testing messages for viruses [30-16](#)
- unclassified [30-7](#)
- virus [30-2](#)
- quarantined messages
  - viewing [30-14](#)
- quarantine level threshold [14-17](#)
- quarantine overflow [14-10](#)
- quarantines
  - centralized policy, virus, and outbreak quarantines [30-10](#)
  - DLP [17-42](#)
  - policy [30-2](#)
  - policy, virus, and outbreak
    - centralized [30-10, 42-5](#)
  - policy, virus, and outbreak, managing [30-3](#)
  - types [30-2](#)
- Quarantine Threat Level Threshold
  - recommended default [14-8](#)
  - setting [14-8](#)
- queries
  - acceptance [25-19](#)
  - chain queries [25-28](#)
  - domain-based [25-26](#)
  - external authentication [25-40](#)
  - group [25-23](#)
  - masquerading [25-21](#)
  - routing [25-20](#)
  - SMTP authentication [25-33](#)
  - spam quarantine alias consolidation [25-44](#)
  - spam quarantine end-user authentication [25-43](#)
- query interface [33-59](#)
- queue [5-3](#)
- quit command [2-9](#)



**R**

RADIUS external authentication [32-22](#)

RAM [40-26](#)

RAM Utilization [34-4](#)

RAT

    bypassing recipients [8-5](#)

    bypassing recipients (CLI) [8-5](#)

    bypassing recipients (GUI) [8-5](#)

rate command [34-17](#)

rate limiting [7-12](#)

rates [34-6](#)

RBL [9-14](#)

RCPT TO [9-10, 9-11, 11-8](#)

RCPT TO command [8-3, 24-7](#)

real-time monitoring [34-16](#)

reboot command [33-2](#)

Received: Header, disabling [5-11](#)

received header [5-11, 13-20](#)

receiving control, bypass [8-5](#)

receiving email, configuring to [5-1](#)

receiving errors [40-23](#)

Recipient Access Table (RAT)

    default entry [8-2](#)

    definition [8-1](#)

    editing via CLI [8-2](#)

recipients, counting in message filters [9-30](#)

recipient validation [22-1](#)

reconfigure [3-13](#)

recursive DNS queries [33-55](#)

recursive entries

    in alias tbales [24-8](#)

    in SMTP Routes [24-2](#)

recursive queries, LDAP [25-14](#)

redirecting email [3-18, 24-2](#)

redirecting URLs in messages [15-9](#)

redirectrecipients [34-26](#)

regional scanning [13-6](#)

regular expressions

DLP [17-15](#)

rejected connections [29-3](#)

relaying email [7-2](#)

relaying messages [3-27, 5-1](#)

remote [33-18](#)

remote upgrades [33-19](#)

removemessage command [34-34](#)

reporting

    DLP [17-42](#)

    Incoming Relays [13-23](#)

reports

    archiving [28-35](#)

reputation filtering

    file [16-1](#)

    sender [6-1](#)

    URL [15-1](#)

required TLS [23-8](#)

resetconfig command [33-4](#)

resetcounters command [28-32, 34-21](#)

resetting [33-4](#)

Resource Conservation mode [34-4, 40-26](#)

resume command [33-3, 34-31](#)

resumedel command [34-29](#)

resumelister command [34-30](#)

resuming email delivery [34-29](#)

resuming receiving [34-30](#)

Retention Time

    for quarantines [30-3](#)

retrospective verdict [16-14](#)

retry message delivery [28-16](#)

reverse DNS [29-5](#)

Reverse DNS Lookup

    disabling [33-56](#)

reverse DNS lookup [7-9, 24-60, 34-19](#)

    timeout [33-54](#)

revert

    installation [33-30](#)

rewriting email addresses [24-7](#)

rewriting URLs in messages

## RFC

- 1035 [24-8](#)
- 1065 [34-36](#)
- 1066 [34-36](#)
- 1067 [34-36](#)
- 1213 [34-36](#)
- 1907 [34-36](#)
- 2047 [16-9, 30-6](#)
- 2487 [23-1](#)
- 2821 [1-11, 5-9](#)
- 821 [10-3](#)
- 822 [10-3](#)

risk factor score [17-2](#)

DLP [17-18](#)

rolling over log files [38-43](#)

rollovernow command [38-7](#)

root servers (DNS) [3-17, 3-26](#)

round robin Virtual Gateways [24-61](#)

routing

SMTP call-ahead server [22-7](#)

routing taking precedence over selected interface [B-3](#)

RSA's DLP Datacenter [17-26, 17-30](#)

RSA Email DLP [17-4](#)

RSA Enterprise Manager [17-23, 30-12](#)

certificates [17-26](#)

clustered appliances [17-32](#)

enabling [17-29](#)

outgoing mail policies [17-31](#)

switching modes [17-33](#)

troubleshooting [31-13](#)

workqueue [31-7](#)

sandboxing. See file analysis

saveconfig command [33-13](#)

SBRS

none [7-7, 9-35](#)

testing [6-6](#)

SBRS score [29-6](#)

SBRS see *Senderbase Reputation Service Score*

scanconfig

scanning recursive levels of attachments [9-116](#)

scanconfig

setting max size for files to be scanned [9-116](#)

skipping attachment types [9-116](#)

scannable archive file types [9-31](#)

scanning images [9-80](#)

Scanning Logs [38-3](#)

scheduled log rollover [38-43](#)

scp command [A-5](#)

SCP Push [38-7](#)

SDS. See Cisco Web Security Services

secure copy [A-5](#)

secure HTTP (https) [23-1](#)

Secure LDAP [25-14](#)

Secure Socket Layer (SSL) [23-1](#)

Security Management appliance [42-1](#)

selecting a notification [21-20](#)

sender

adding senders to sender groups via GUI [7-14](#)

SenderBase [5-14, 7-11, 7-17, D-1](#)

SBO in sender groups [7-7](#)

timeout per connection [5-11](#)

using IP Profiling [5-11](#)

SenderBase, querying [7-7](#)

SenderBase Affiliate network [6-1](#)

SenderBase Network Owner Identification Number [7-4](#)

SenderBase Reputation Score [6-2, 7-7, 7-13, 40-2](#)

SenderBase Reputation score [13-22](#)

SenderBase reputation score [29-6](#)

---

**S**

safelist/blocklist [31-7](#)

and external spam quarantine [31-8](#)

backing up and restoring [31-13](#)

enabling [31-8](#)

importing and exporting [31-13](#)

managing [31-9](#)

on multiple ESAs [31-12](#)

- SenderBase Reputation Scores, syntax in CLI [7-7](#)
- SenderBase Reputation Service [6-1, 28-1, 28-13](#)
- SenderBase Reputation Service Score [7-6](#)
- sender group
  - adding via GUI [7-13](#)
  - BLACKLIST [7-11](#)
  - SUSPECTLIST [7-11](#)
  - UNKNOWNLIST [7-12](#)
  - WHITELIST [7-11](#)
- sender groups
  - overview [7-3](#)
- sender rate limit
  - exceeded error code [7-17](#)
  - exceeded error text [7-17](#)
  - exceptions [7-17](#)
  - maximum recipients per time interval [7-17](#)
- sender verification
  - malformed MAIL FROM and default domain [7-30](#)
- sender verification exception table [7-30](#)
- separate window icon [28-7](#)
- serial connection pinouts [A-5](#)
- serv.fail [7-38](#)
- SERVFAIL [7-29, 7-38](#)
- services for interfaces [A-1](#)
- sethostname command [33-53](#)
- setup [3-1](#)
- severity level [17-21](#)
- severity scale [17-21](#)
  - DLP [17-21](#)
- severity settings [17-21](#)
- showconfig command [33-12](#)
- showmessage command [34-34](#)
- showrecipients [34-27](#)
- shutdown command [33-2](#)
- SIDF records
  - testing [20-23](#)
  - valid [20-23](#)
- SIDF verification [9-11](#)
  - configuring [20-22](#)
  - conformance level [20-24](#)
  - enabling [20-24](#)
  - results [20-31](#)
  - testing [20-34](#)
- significant bits
  - set in mail flow policy [7-17](#)
- signing
  - DKIM [20-2](#)
  - Domain Keys [20-2](#)
  - dual Domain Keys and DKIM [20-2](#)
- signing key
  - size [20-3](#)
- signing keys
  - deleting all existing keys [20-13](#)
  - removing specific keys [20-12](#)
- SMI file [34-37](#)
- SMTP [D-1](#)
  - banner host name [7-16](#)
  - banner text [7-8](#)
  - code [7-8](#)
  - HELO command [7-11](#)
  - messages [5-15](#)
  - response [8-3](#)
  - testing IronPort Anti-Spam [13-25](#)
- SMTP Address Parsing
  - Loose mode [5-9](#)
  - Strict mode [5-9](#)
- SMTP Auth [25-2, 25-32](#)
  - DIGEST-MD5 [25-36](#)
  - MD5 [25-33](#)
  - SHA [25-33](#)
  - supported authentication mechanisms [25-33](#)
  - TLS [25-37](#)
- SMTP authenticated user match filter rule [9-40](#)
- SMTP Authentication [29-5](#)
  - HAT entry [7-19](#)
- SMTP Authentication profile [25-36](#)
- SMTP Authentication with forwarding
  - defined [25-35](#)

- SMTP call-ahead recipient validation [22-1](#)
  - bypassing [22-8](#)
  - conversation workflow [22-2](#)
  - SMTP server responses [22-5](#)
  - with LDAP routing query [22-6](#)
- SMTP CALL-Ahead Server Profile
  - creating [22-3](#)
  - enabling on a listener [22-6](#)
- SMTP conversation
  - SMTP call-ahead server [22-2](#)
- SMTP daemon
  - see *injector*
  - see *listener*
- SMTP HELO command [40-22](#)
- SMTP query workflow [22-7](#)
- SMTP Routes [24-1](#)
  - limits [24-3](#)
  - mail delivery and splintering [24-4](#)
  - multiple host entries [24-3](#)
  - recursive entries in [24-2](#)
  - USEDNS [24-3](#)
- SMTP Routes, maximum [24-2](#)
- SMTP Routes and DNS [24-3](#)
- SNMP
  - community string [34-37](#)
  - hardware failure trap conditions [34-38](#)
  - IPMI [34-37](#)
  - MIB file [34-37](#)
  - overview [34-36](#)
  - SMI file [34-37](#)
  - specifying multiple trap targets [34-39](#)
  - traps [34-39](#)
- SNMP (Simple Network Management Protocol) [34-36](#)
- SNMPv1 [34-37](#)
- SNMPv2 [34-37](#)
- Sophos
  - evaluation key [3-21, 3-34, 12-2](#)
  - updates [12-18](#)
- Sophos virus scanning
  - filters [12-11](#)
- source routing [5-10](#)
- spam
  - altering the subject line of [13-9](#)
  - archiving [13-9](#)
  - including a custom header in [13-9](#)
  - sending to an alternate address [13-9](#)
  - sending to an alternate mailhost [13-9](#)
  - testing [13-25](#)
- spam message [28-8](#)
- spam quarantine
  - alias consolidation [31-21](#)
  - behavior when full [31-3](#)
  - deleting all messages [31-24](#)
  - disabling [31-24](#)
  - end user access [31-16](#)
  - end-user access [31-1, 31-14](#)
  - external [31-1, 42-3](#)
  - IMAP/POP authentication [31-16](#)
  - LDAP authentication [31-15](#)
  - local [31-1](#)
  - message details [31-23](#)
  - message variables [31-19](#)
  - notification [31-19](#)
  - receiving multiple notifications [31-20](#)
  - released messages and email pipeline [31-24](#)
  - safelist/blocklist. See safelist/blocklist.
  - testing notifications [31-21](#)
- specifying an offset [33-59](#)
- spf-passed filter rule [9-11, 20-33](#)
- SPF records
  - testing [20-23](#)
  - valid [20-22](#)
- spf-status filter rule [9-11, 20-32](#)
- SPF verification
  - configuring [20-22](#)
  - conformance level [20-24](#)
  - enabling [20-24](#)
  - received SPF header [20-30](#)

- results [20-31](#)
- testing [20-34](#)
- SPFverification [9-11](#)
- square brackets [2-4](#)
- SSH [2-3, D-1](#)
- SSL [25-14](#)
- STARTTLS
  - definition [23-1](#)
- stateless logs [38-15](#)
- static routes [24-56](#)
- status command [34-7](#)
- status detail command [34-8](#)
- Status Logs [38-2](#)
- stopped by content filter [28-8](#)
- stopped by reputation filtering [28-8](#)
- streaming upgrades [33-19](#)
- strip-header filter action [9-69](#)
- strip headers [9-69](#)
- stripping subdomains [24-16](#)
- subject "No Subject" [29-5](#)
- subnet [3-18, 3-26](#)
- supported languages
  - configuring default [33-61](#)
- SUSPECTLIST sender group [7-11](#)
- suspend command [33-2](#)
- suspenddel command [34-28](#)
- suspending email delivery [34-28](#)
- suspending receiving [34-29](#)
- suspendlistener command [34-30](#)
- suspicious senders, throttling [7-11](#)
- synchronizing time [3-16, 3-36](#)
- Syslog [38-7](#)
- System Capacity
  - All page [28-30](#)
  - Incoming Mail page [28-26](#)
  - memory page swapping [28-29](#)
  - Outgoing Mail page [28-27](#)
  - System Load page [28-28](#)
  - WorkQueue page [28-25](#)

- System Capacity page [28-24](#)
- system clock [3-16, 3-36](#)
- System Logs [38-2](#)
- system monitoring through the GUI [36-7](#)
- system quarantine. See quarantines, policy, virus, and outbreak
- system setup [3-1](#)
- systemsetup command [3-24](#)
- system setup next steps [3-23](#)
- system setup wizard [3-13](#)
- System Status page [28-30](#)
- system time
  - setting [3-16, 3-36](#)

---

## T

- tail command [38-45](#)
  - parameters [38-45](#)
- TCP listen queue [5-11](#)
- TCPREFUSE [7-8](#)
- Telnet [2-3, A-1, D-1](#)
- testing
  - IronPort Anti-Spam [13-25](#)
  - Sophos virus engine [12-16](#)
  - system setup [3-36](#)
- testing HAT variables [7-10](#)
- text resources
  - code view [21-11](#)
  - content dictionary [21-1](#)
  - disclaimers [21-12](#)
  - exporting [21-10](#)
  - exporting and importing into HTML resources [21-11](#)
  - HTML-based [21-11](#)
  - importing [21-10](#)
  - managing [21-9](#)
  - non-ASCII characters [21-8](#)
  - understanding [21-8](#)
  - using in policies and settings [21-12](#)
- third-party relay [8-2](#)

Threat Level  
     defined [14-6](#)

Threat Operations Center (TOC) [14-6, 28-6](#)

thresholds, in SenderBase Reputation Scores [7-7](#)

throttling [6-1, 7-11](#)

time, system [3-16, 3-36](#)

time servers [3-16, 3-36](#)

time zone [33-59](#)

time zone, setting [3-16, 3-36](#)

time zone files  
     updating [33-59](#)

Time Zone page [33-59](#)

TLS  
     certificates [23-1](#)  
     default [23-10](#)  
     preferred [23-10](#)  
     required [23-11](#)

tlsverify command [40-25](#)

CASE (Context Adaptive Scanning Engine) [13-23](#)

tophosts command [34-15, 40-20](#)

topin command [34-19](#)

trace [13-23](#)

trace command [6-6, 40-1](#)

Trace page [40-1](#)

tracking  
     "AND" searches [29-2](#)

Transport Layer Security (TLS) [7-18](#)

troubleshooting  
     DLP [17-42](#)

troubleshooting delivery [40-23](#)

trustworthiness [7-7](#)

TTL [34-12](#)

tzupdate  
     CLI command [33-59](#)

UNKNOWNLIST sender group [7-12](#)

unparsable messages [9-23](#)

unsolicited commercial email [6-1](#)

updates  
     DLP engine and classifiers [17-39](#)

update server [33-22](#)

upgrades [33-18](#)  
     available [33-27, 33-29](#)  
     local [33-18](#)  
     obtaining via GUI [33-19](#)  
     remote [33-19](#)  
     streaming [33-18, 33-19](#)

upgrade server [33-19](#)

upstream proxy  
     file reputation [16-6](#)

URL list. See whitelist, URL filtering.

URL reputation [15-1](#)

user accounts [32-1](#)  
     limits [32-1](#)  
     locking and unlocking [32-17](#)

user groups [32-1, 32-2](#)

user name [32-4](#)

user preferences  
     defining [33-61](#)

user types [32-2](#)

using HAT variables [7-9](#)

uuencoded attachments [9-6](#)

---

## U

unary form, in message filters [9-30](#)

unclassified quarantine. See quarantine, unclassified

---

## V

verdict  
     image analysis [11-6, 11-11](#)

verification  
     SIDF [20-22](#)  
     SPF [20-22](#)

verifying senders  
     exception table [7-35](#)

version [28-31](#)

virtual appliance

license [38-2](#)

Virtual Domains [24-16](#)

virtual Email Security appliance

- loading the license [3-8](#)

Virtual Gateway™ technology [24-59](#)

Virtual Gateway addresses [9-68, 24-63](#)

Virtual Gateway queue [24-60](#)

Virtual IP (VIP) [37-13](#)

virtual table [24-28](#)

virus definition

- automatic update interval [33-23](#)

virus message [28-8](#)

virus quarantine. See quarantine

- virus.

Virus Types page [28-21](#)

virususerstable. See alias tables

VLAN

- defined [37-6](#)
- labels [37-7](#)

VRT. See file analysis.

Active Directory [3-22](#)

- system setup [3-1, 3-13](#)

work queue [34-5, 34-32](#)

work queue, pausing [34-32](#)

---

## X

X.509 certificate [23-2](#)

X-advertisement header [13-25](#)

X-headers, adding [16-9, 30-6](#)

X-IronPort-Anti-Spam header [13-14](#)

X-IronPort-AV header [12-8](#)

XML [33-7, 33-8, 33-10, 33-13, 36-11, 38-2](#)

XML Status feature [36-11](#)

---

## W

WBRS

- See URL reputation

web interface

- enabling [3-26](#)

web reputation

- message filters [9-47, 9-75](#)

Web UI session timeout [32-26](#)

weekly status updates [3-35](#)

whitelist

- URL filtering [15-4](#)

WHITELIST sender group [7-11, 12-13](#)

white space [9-17](#)

whitespace [12-10, 13-9](#)

whoami command [32-6](#)

who command [32-6](#)

wizard