



Managing Spam and Graymail

This chapter contains the following sections:

- [Overview of Anti-Spam Scanning](#) , on page 1
- [How to Configure the Appliance to Scan Messages for Spam](#), on page 2
- [IronPort Anti-Spam Filtering](#), on page 3
- [Configuring Intelligent Multi-Scan and Graymail Detection](#), on page 7
- [Defining Anti-Spam Policies](#) , on page 18
- [Protecting Appliance -Generated Messages From the Spam Filter](#), on page 24
- [Headers Added During Anti-Spam Scanning](#) , on page 25
- [Reporting Incorrectly Classified Messages to Cisco](#), on page 25
- [Determining Sender IP Address In Deployments with Incoming Relays](#) , on page 30
- [Monitoring Rules Updates](#), on page 39
- [Testing Anti-Spam](#), on page 40

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
 - 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](#).

Related Topics

- [Anti-Spam Solutions](#) , on page 2

Anti-Spam Solutions

Your appliance offers the following anti-spam solutions:

- [IronPort Anti-Spam Filtering](#), on page 3.
- [Configuring Intelligent Multi-Scan and Graymail Detection](#), on page 7.

You can license and enable both these solutions on your 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

Procedure

	Command or Action	Purpose
Step 1	Enable anti-spam scanning on the appliance .	<p>Note Remaining steps in this table apply to both scanning engine options.</p> <p>If you have feature keys for both Cisco IronPort Anti-Spam and Intelligent Multi-Scan, you can enable both solutions on the appliance .</p> <ul style="list-style-type: none"> • IronPort Anti-Spam Filtering, on page 3 • Configuring Intelligent Multi-Scan and Graymail Detection, on page 7
Step 2	Configure whether to quarantine spam on the local appliance or use an external quarantine on a Security Management appliance .	<ul style="list-style-type: none"> • Setting Up the Local Spam Quarantine • Working with an External Spam Quarantine
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

	Command or Action	Purpose
Step 4	Configure the anti-spam scanning rules for the user groups you defined.	Defining Anti-Spam Policies , on page 18
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
Step 6	(Recommended) Enable IP Reputation Service scoring for each inbound mail flow policy, even if you are not rejecting connections based on IP 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 .
Step 7	If your 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 , on page 30
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 , on page 24
Step 9	(Optional) Enable URL filtering to strengthen protection against malicious URLs in messages.	Enable URL Filtering
Step 10	Test your configuration.	Testing Anti-Spam , on page 40
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. <ul style="list-style-type: none"> • Service Updates • Updates Through a Proxy Server • Configuring Server Settings for Downloading Upgrades and Updates

IronPort Anti-Spam Filtering

Related Topics

- [Evaluation Key](#), on page 4
- [Cisco Anti-Spam: an Overview](#) , on page 4
- [Configuring IronPort Anti-Spam Scanning](#), on page 5

Evaluation Key

Your 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 `antispmconfig` 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](#)) 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](#).)

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](#) , on page 4
- [URL-Related Protections and Controls](#)

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, on page 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](#) , on [page 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:
- a) Click **Enable**.
 - b) 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
<p>Message Scanning Thresholds</p>	<p>a. Enter a value for <i>Always scan messages smaller than</i> —The recommended value is 1 MB 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.</p> <p>Cisco advises not to exceed 3 MB for the <i>always scan</i> message size. A larger value may result in decreased performance.</p> <p>b. Enter a value for <i>Never scan messages larger than</i> —The recommended value is 2 MB 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.</p> <p>Cisco advises not to exceed 10 MB for the <i>never scan</i> message size. A larger value may result in decreased performance.</p> <p>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.</p> <p>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.</p>
<p>Timeout for Scanning Single Message</p>	<p>Enter the number of seconds to wait for timeout when scanning a message.</p> <p>Enter an integer from 1 to 120. The default value is 60 seconds.</p>
<p>Scanning Profile</p>	<p>Choose from any of the following scanning profiles to catch spam messages:</p> <ul style="list-style-type: none"> • Normal - Enable this option for a balanced approach to block spam. • Aggressive - Enable this option to provide stronger emphasis to block spam. When enabled, tuning the Anti-Spam policy thresholds have more impact on spam detection than the Normal profile with a larger potential for false positives. <p>Note When using the new aggressive scanning profile mail policy adjustments to Anti-Spam thresholds have a larger impact than before. Therefore when enabling the aggressive profile, any Anti-Spam policy thresholds previously adjusted should be reset to default settings and then reevaluated for the best balance of spam catch rate vs. false positive potential.</p> <ul style="list-style-type: none"> • Regional (China) - Enable this only if you receive the bulk of your email from the specified region. The supported region is China. As this option 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.

Configuring Intelligent Multi-Scan and Graymail Detection

This section describes how to configure Cisco Intelligent Multi-Scan and Graymail Detection and Safe Unsubscribing:

- [Configuring Cisco Intelligent Multi-Scan, on page 7](#)
- [Configuring Global Settings for Intelligent Multi-Scan and Graymail Detection, on page 17](#)
- [Managing Graymail, on page 8](#)

Configuring Cisco Intelligent Multi-Scan

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 Cisco 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.



Important 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](#). You will see the IronPort Intelligent Multi-Scan option only if you have done so.

Procedure

- Step 1** Select **Security Services > IMS and Graymail**.
- 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 IMS Settings**.
- Step 4** Select the check box for **Enable Intelligent Multi-Scan** enable the feature globally for the appliance . However, you must still enable per-recipient settings in Mail Policies.
- Step 5** (Optional) Click **Edit Global Settings** to configure the threshold for message scanning. For more information about global settings, see [Configuring Global Settings for Intelligent Multi-Scan and Graymail Detection, on page 17](#).
- Step 6** Submit and commit your changes.
-

Managing Graymail

- [Overview of Graymail, on page 8](#)
- [Graymail Management Solution in Email Security Appliance , on page 8](#)
- [How Graymail Management Solution Works, on page 9](#)
- [Configuring Graymail Detection and Safe Unsubscribing, on page 12](#)
- [Troubleshooting Graymail Detection and Safe Unsubscribing, on page 17](#)

Overview of Graymail

Graymail messages are messages that do not fit the definition of spam, for example, newsletters, mailing list subscriptions, social media notifications, and so on. These messages were of use at some point in time, but have subsequently diminished in value to the point where the end user no longer wants to receive them.

The difference between graymail and spam is that the end user intentionally provided an email address at some point (for example, the end user subscribed to a newsletter on an e-commerce website or provided contact details to an organization during a conference) as opposed to spam, messages that the end user did not sign up for.

Graymail Management Solution in Email Security Appliance

The graymail management solution in the appliance comprises of two components: an integrated graymail scanning engine and a cloud-based Unsubscribe Service.

The graymail management solution allows organizations to:

- Identify graymail using the integrated graymail engine and apply appropriate policy controls.
- Provide an easy mechanism for end users to unsubscribe from unwanted messages using Unsubscribe Service.

In addition to these, the graymail management solution also help organizations to provide:

- **Secure unsubscribe option for end users.** Mimicking an unsubscribe option is a popular phishing technique. For this reason, the end users are generally wary of clicking unknown unsubscribe links. For such scenarios, the cloud-based Unsubscribe Service extracts the original unsubscribe URI, checks the reputation of the URI, and then performs the unsubscribe process on behalf of the end user. This protects end users from malicious threats masquerading as unsubscribe links.
- **Uniform subscription management interface for end users.** Different graymail senders use different layouts for displaying unsubscribe links to the users. The users must search for the unsubscribe link in the message body and perform the unsubscribing. Irrespective of the graymail senders, the graymail management solution provides a common layout for displaying unsubscribe links to the users.
- **Better visibility for administrators into various graymail categories.** The graymail engine classifies each graymail into three categories (see [Graymail Classification, on page 9](#)) and the administrators can set policy controls based on these categories.
- **Improved spam efficacy**

Related Topics

- [Graymail Classification, on page 9](#)

Graymail Classification


The graymail engine classifies each graymail into one of the following categories:

- **Marketing Email.** Advertising messages sent by professional marketing groups, for example, bulletins from Amazon.com with details about their newly launched products.
- **Social Network Email.** Notification messages from social networks, dating websites, forums, and so on. Examples include alerts from:
 - LinkedIn, for jobs that you may be interested in
 - CNET forums, when a user responds to your post.
- **Bulk Email.** Advertising messages sent by unrecognized marketing groups, for example, newsletters from TechTarget, a technology media company.

How Graymail Management Solution Works

The following steps illustrates the workflow of graymail management solution:

Figure 1: Graymail Management Solution Workflow



Workflow

Procedure

-
- Step 1** The appliance receives an incoming message.
 - Step 2** The appliance checks if graymail detection is enabled. If graymail detection is enabled, go to Step 3. Else, go to Step 8
 - Step 3** The appliance checks if the message is spam, virus, or malware positive. If positive, go to Step 8. Else, go to Step 4
 - Step 4** The appliance checks if the message is graymail. If the message is graymail, go to Step 5. Else, go to Step 8
 - Step 5** The appliance applies the configured policy actions such as, drop, deliver, bounce, or quarantine to the spam quarantine.
 - Step 6** The appliance checks if safe unsubscribing enabled. If safe unsubscribing is enabled, go to Step 7. Else, go to Step 8.
 - Step 7** The appliance adds a banner with unsubscribe button to the message. Also, the appliance rewrites the existing unsubscribe links in the message body.
 - Step 8** The appliance processes the message through the next stages of its email work queue.
-

What to do next

For an overview of how email is processed through the system, from reception to routing to delivery, see [Understanding the Email Pipeline](#)

Related Topics


- [How Safe Unsubscribing Works, on page 11](#)

- Understanding the Email Pipeline

How Safe Unsubscribing Works

The following flow diagram shows how safe unsubscribing works.

Figure 2: Safe Unsubscribing Workflow



Workflow

Procedure

-
- Step 1** End user receives a message with the graymail banner.
 - Step 2** End user clicks on the Unsubscribe link.

- Step 3** Unsubscribe Service extracts the original unsubscribe URI.
- Step 4** Unsubscribe Service checks the reputation of the URI.
- Step 5** Depending on the reputation of the URI, the Unsubscribe Service performs one of the following actions:
- If the URI is malicious, the Unsubscribe Service will not perform the unsubscribe process and displays a block page to the end user.
 - If the URI is not malicious, depending on the URI type (http or mailto), the Unsubscribe Service sends an unsubscribe request to the graymail sender.
 - If the request is successful, the Unsubscribe Service displays the “Successfully unsubscribed” status to the end user.
 - If the first unsubscribe request fails, the Unsubscribe Service displays the “Unsubscribe process in progress” status and provides a URL that can be used to track the status of the unsubscribing.

End users can use this URL to track the status at a later point. After the first failed attempt, the Unsubscribe Service sends periodic unsubscribe requests for a duration of four hours.

If an end user checks the status of the unsubscribe process at a later point,
 - If one of the requests within the four hour duration (from the first failed attempt) is successful, the Unsubscribe Service displays the “Successfully unsubscribed” status to the end user.
 - If none of the requests within the four hour duration (from the first failed attempt) are successful, the Unsubscribe Service displays the “Unable to subscribe” status to the end user and provides a URL that can be used to unsubscribe from the graymail manually.

Configuring Graymail Detection and Safe Unsubscribing

- [Requirements for Graymail Detection and Safe Unsubscribing, on page 12](#)
- [Graymail Detection and Safe Unsubscribing in Cluster Configurations, on page 13](#)
- [Enable Graymail Detection and Safe Unsubscribing, on page 13](#)
- [Configuring the Incoming Mail Policy for Graymail Detection and Safe Unsubscribing, on page 13](#)
- [IronPort-PHdr Header Added During Graymail Scanning, on page 14](#)
- [Bypassing Graymail Actions using Message Filters, on page 15](#)
- [Monitoring Graymail, on page 15](#)
- [Updating Graymail Rules, on page 16](#)
- [Customizing the Appearance of Unsubscribe Page for End Users, on page 16](#)
- [End-User Safelist, on page 16](#)
- [Viewing Logs, on page 16](#)

Requirements for Graymail Detection and Safe Unsubscribing

- For graymail detection, anti-spam scanning must be enabled globally. This can be either the IronPort Anti-Spam, the Intelligent Multi-Scan feature, or Outbreak Filters. See [Managing Spam and Graymail, on page 1](#).
- For safe unsubscribing,
 - Add the safe unsubscribing feature key.

- The end user machines must be able to connect to the cloud-based Unsubscribe Service directly over the Internet.

Graymail Detection and Safe Unsubscribing in Cluster Configurations

You can enable Graymail Detection and Safe Unsubscribing at the machine, group or cluster level.

Enable Graymail Detection and Safe Unsubscribing

Procedure

- Step 1** Select **Security Services > IMS and Graymail**.
- Step 2** Click **Edit Graymail Settings**.
- Step 3** Check **Enable Graymail Detection**.
- Step 4** Check **Enable Safe Unsubscribe**.
- Step 5** (Optional) Check **Enable Automatic Updates** to enable automatic update of the engine.
The appliance fetches the required updates for the particular engine from the update server.
- Step 6** Click **Submit**.
- Step 7** (Optional) Click **Edit Global Settings** to configure the threshold for scanning the message. For more information, see [Configuring Global Settings for Intelligent Multi-Scan and Graymail Detection, on page 17](#).
- Step 8** Submit and commit your changes.
-

What to do next

To configure Graymail Detection and Safe Unsubscribing global settings in CLI, use the `imsandgraymailconfig` CLI command. For more information, see *CLI Reference Guide for AsyncOS for Cisco Email Security Appliances*.

Configuring the Incoming Mail Policy for Graymail Detection and Safe Unsubscribing

Before You Begin

[Enable Graymail Detection and Safe Unsubscribing, on page 13](#)

Procedure

- Step 1** Click **Mail Policies > Incoming Mail Policies**.
- Step 2** Click the link in the **Graymail** column of the mail policy to modify.
- Step 3** Depending on your requirements, choose the following options:
- Enable graymail detection
 - Enable safe unsubscribing
 - Choose whether to apply the above actions on all messages or only on unsigned messages.

Note The appliance considers a message signed if it is encrypted using S/MIME or it contains an S/MIME signature.

- Actions to be taken on various graymail categories (Marketing Email, Social Network Email, and Bulk Email):

- Drop, deliver, bounce, or quarantine (to the spam quarantine) the message

Note If you plan to use safe unsubscribing option, you must set the action to deliver or quarantine.

- Send the message to an alternate host
- Modify subject of the message
- Add custom headers
- Send the message to an alternate envelope recipient

Note If you are sending a graymail positive message to an alternate envelope recipient, banner will not be added.

- Archive the message

Note If you are planning only to monitor the detected graymail, you can enable graymail detection per policy without having to configure actions for various graymail categories. In this scenario, the appliance takes no action on the detected graymail.

Step 4 Submit and commit your changes.

What to do next



Note You can also configure outgoing mail policies for graymail detection. Keep in mind that, in this scenario, you cannot configure safe unsubscribing.

To configure policy settings for Graymail Detection and Safe Unsubscribing in CLI, use the **policyconfig** command. For more information, see *CLI Reference Guide for AsyncOS for Cisco Email Security Appliances*.

IronPort-PHdr Header Added During Graymail Scanning

The IronPort-PHdr header is added to all messages that are processed by the Graymail engine when:

- Graymail engine is enabled globally on the appliance .
- Graymail scanning is enabled for a specific mail policy.



Note If Graymail scanning is not enabled for a specific mail policy, the IronPort-PHdr header is still added to all messages, if the Graymail engine is enabled globally on the appliance .

The IronPort-PHdr header contains encoded proprietary information and is not customer-decodable. This header provides additional information about debugging issues with your Graymail configuration.



Note If Anti-Spam engine or Outbreak Filter is enabled for a specific mail policy, the IronPort-PHdr header is added to all messages that pass through the specific mail policy.

Bypassing Graymail Actions using Message Filters

If you do not want to apply graymail actions on certain messages, you can use the following message filters to bypass graymail actions:

Message Filter Action	Description
skip-marketingcheck	Bypass actions on marketing emails
skip-socialcheck	Bypass actions on social network emails
skip-bulkcheck	Bypass actions on bulk emails

The following example specifies that messages received on the listener “private_listener” must bypass graymail actions on social network emails.

```
internal_mail_is_safe:
if (recv-listener == 'private_listener')
{
skip-socialcheck
();
}
```

Monitoring Graymail

You can view data about detected graymail using the following reports.

Report	Contains the Following Graymail Data	More Info
Overview page > Incoming Mail Summary	The number of incoming graymail messages under each graymail category (Marketing, Social, and Bulk) and the total number of graymail messages.	Overview Page
Incoming Mail page > Top Senders by Graymail Messages	The top graymail senders.	Incoming Mail Page
Incoming Mail page > Incoming Mail Details	The number of incoming graymail messages under each graymail category (Marketing, Social, and Bulk) and the total number of graymail messages for all the IP addresses, domain names, or network owners.	
Incoming Mail page > Incoming Mail Details > Sender Profile (drill down view)	The number of incoming graymail messages under each graymail category (Marketing, Social, and Bulk) and the total number of graymail messages for a given IP address, domain name, or network owner.	

Report	Contains the Following Graymail Data	More Info
Internal Users page > Top Users by Graymail	The top end users who receive graymail.	Internal Users Page
Internal Users page > User Mail Flow Details	The number of incoming graymail messages under each graymail category (Marketing, Social, and Bulk) and the total number of graymail messages for all the users.	
Internal Users page > User Mail Flow Details > Internal User (drill down view)	The number of incoming graymail messages under each graymail category (Marketing, Social, and Bulk) and the total number of graymail messages for a given user.	

If you had enabled Marketing Email Scanning under anti-spam settings for a mail policy, after upgrading to AsyncOS 9.5 or later, keep in mind that:

- The number of marketing messages is a sum of marketing messages detected before and after the upgrade.
- The total number of graymail messages does not include the number of marketing messages detected before the upgrade.
- The total number of attempted messages also includes the number of marketing messages detected before the upgrade.

Updating Graymail Rules

If you have enabled service updates, scanning rules for the graymail management solution is retrieved from the Cisco update servers. But in some scenarios (for example, you have disabled automatic service updates or automatic service update is not working), you may want to manually update graymail rules.

To manually update the graymail rules, do one of the following:

- In web interface, go to **Security Service > IMS and Graymail** page, and click **Update Now**.
- In CLI, run the `graymailupdate` command.

To know the details of existing graymail rules, see the **Rule Updates** section of the **IMS and Graymail** page in web interface or use the `graymailstatus` command in CLI.

Customizing the Appearance of Unsubscribe Page for End Users

When an end user clicks on unsubscribe link, the Unsubscribe Service displays a Cisco branded Unsubscribe page indicating the status of the unsubscribe process (see [How Safe Unsubscribing Works, on page 11](#)). You can customize the appearance of the Unsubscribe page and display your organization's branding (such as company logo, contact information, and so on) using **Security Services > Block Page Customization**. For instructions, see [Customizing the Notification That End Users See If a Site Is Malicious](#).

End-User Safelist

If the end users in your organization have configured Safelist for their own email accounts, graymail messages from a sender in the safelist will not be scanned by the graymail scanning engine. For more information about Safelists, see [Using Safelists and Blocklists to Control Email Delivery Based on Sender](#).

Viewing Logs

The graymail detection and safe unsubscribing information is posted to the following logs:

- **Graymail Engine Logs**. Contains information about the graymail engine, status, configuration, and so on. Most information is at Info or Debug level.

- **Graymail Archive.** Contains archived messages (the messages that are scanned and associated with the “archive message” action). The format is an mbox-format log file.
- **Mail Logs.** Contains information about graymail detection and addition of banner for safe unsubscribing. Most information is at Info or Debug level.

Troubleshooting Graymail Detection and Safe Unsubscribing

[Unable to Perform Safe Unsubscribing, on page 17](#)

Unable to Perform Safe Unsubscribing

Problem

After clicking on the Unsubscribe link, the end user sees the following message: “Unable to unsubscribe from...”

Solution

This problem can occur if the Unsubscribe Service is unable to perform the safe unsubscribe on behalf of the end user. The following are some of the common scenarios in which the Unsubscribe Service is unable to perform the safe unsubscribe:

- Unsubscribe URI or mailto address is wrong.
- Websites that require the end users’ credentials to unsubscribe.
- Websites that require the end users to confirm the request of unsubscribing by logging into their email accounts.
- Websites that require captcha to be solved and the Unsubscribe Service is unable to solve the captcha.
- Websites that require interactive unsubscribing.

The end users can use the URL provided at the bottom of the unsubscribe page to unsubscribe manually.

Configuring Global Settings for Intelligent Multi-Scan and Graymail Detection

To optimize the throughput of your appliance, you can configure the threshold and the timeout settings for scanning messages by Cisco Intelligent Multi-Scan and Graymail. These settings are common for both Cisco Intelligent Multi-Scan and Graymail configuration.

1. Select **Security Services > IMS and Graymail**
2. Click **Edit Global Settings**.
3. Select the thresholds for scanning with Cisco Intelligent Multi-Scan and Graymail Detection.

The default values are:

- Always scan 512K or less.



Note This setting is not applicable for Graymail Detection and Safe Unsubscribing.

- Never scan 1M or more.

4. 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 do not have 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.

5. 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, on page 2](#).
- Familiarize yourself with the following:
 - [Understanding Positive and Suspect Spam Thresholds, on page 20](#)
 - [Configuration Examples: Actions for Positively Identified versus Suspected Spam , on page 21](#)
 - [Unwanted Marketing Messages From Legitimate Sources, on page 21](#)
 - If you have enabled more than one anti-spam solution: [Enabling Different Anti-Spam Scanning Engines in Different Mail Policies: Configuration Example , on page 23](#)
 - [Headers Added During Anti-Spam Scanning , on page 25](#)
- If you will archive spam into the “Anti-Spam Archive” log, see also [Logging](#).
- If you will send messages to an alternate mailhost, see also [Alter Delivery Host Action](#).

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 Enable Marketing Email Scanning	Choose an option. 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: <ul style="list-style-type: none"> • Deliver • Drop • Bounce • Quarantine
(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 .
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 not 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 <code>[SPAM]</code> with a few trailing spaces. “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 , on page 22.

Option	Description
(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](#).

Related Topics

- [How to Configure the Appliance to Scan Messages for Spam, on page 2](#)
- [Understanding Positive and Suspect Spam Thresholds, on page 20](#)
- [Configuration Examples: Actions for Positively Identified versus Suspected Spam , on page 21](#)
- [Unwanted Marketing Messages From Legitimate Sources, on page 21](#)
- [Using Custom Headers to Redirect URLs in Suspected Spam to the Cisco Web Security Proxy: Configuration Example , on page 22](#)
- [Enabling Different Anti-Spam Scanning Engines in Different Mail Policies: Configuration Example , on page 23](#)

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](#) , on page 2
- [Configuration Examples: Actions for Positively Identified versus Suspected Spam](#) , on page 21

Configuration Examples: Actions for Positively Identified versus Suspected Spam

Spam	Sample Actions (Aggressive)	Sample Actions (Conservative)
Positively Identified	Drop	<ul style="list-style-type: none"> • 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 [Managed Exceptions](#).

Unwanted Marketing Messages From Legitimate Sources

If you had configured Marketing Email Settings under anti-spam settings for a mail policy, after upgrading to AsyncOS 9.5 for Email, Marketing Email Settings under anti-spam settings will be moved under graymail settings of the same policy. See [Managing Spam and Graymail, on page 1](#).

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](#).

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.
 - In the Suspected Spam Settings section, enable suspected spam scanning.
 - Click **Advanced** to display the Add Custom Header option.
 - Add a custom header such as `url_redirect`.
 - Submit and commit your changes.
- Step 2** Create a content filter to redirect URLs in messages that have the custom header:
- Select **Mail Policies > Incoming Content Filters**.
 - Click **Add Filter**.
 - Name the filter `url_redirect`.
 - Click **Add Condition**.
 - Click **Other Header**.
 - Enter the header name: `url_redirect`.
- Make sure this exactly matches the header you created above.
- Select **Header exists**.
 - Click **OK**.
 - Click **Add Action**.
 - Click **URL Category**.
 - Select all categories in **Available Categories** and add them to **Selected Categories**.
 - For Action on URL, select **Redirect to Cisco Security Proxy**.
 - Click **OK**.
- Step 3** Add the content filter to the mail policy.
- Select **Mail Policies > Incoming Mail Policies**.
 - Click the link in the **Content Filters** column for the policy that you selected earlier in this procedure.
 - Select **Enable Content Filters** if it is not already selected.
 - Select the check box to enable the **url_filtering** content filter.
 - Submit and commit your changes.
-

What to do next**Related Topics**

- [Redirecting URLs](#)
- [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 3: Mail Policies - Anti -spam Engine Per Recipient

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	(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.

The following figure shows Cisco Intelligent Multi-Scan and unwanted marketing message detection enabled in a policy.

Figure 4: Mail Policies - Enabling Cisco Intelligent Multi-scan

After submitting and committing the changes, the mail policy looks like this:

Figure 5: Mail Policies - Intelligent Multi-Scan Enabled in Policy

Incoming Mail Policies

Find Policies

Email Address: Recipient Sender

Policies

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)	<input type="button" value="Delete"/>
	Default Policy	IronPort Anti-Spam Positive: Deliver Suspected: Deliver Marketing Messages: Disabled	Not Available	Disabled	Not Available	

Key:

Protecting Appliance -Generated Messages From the Spam Filter

Because automated email messages that are sent from the 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](#) and [Bypass Anti-Spam System Action](#).

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](#) , on page 18.

Related Topics

- [Using Custom Headers to Redirect URLs in Suspected Spam to the Cisco Web Security Proxy: Configuration Example](#) , on page 22

Reporting Incorrectly Classified Messages to Cisco

Messages that appear to be incorrectly classified may be reported to Cisco for analysis. The reported messages are used to enhance the accuracy and effectiveness of the product.

You can report incorrectly classified messages that belong to the following categories:

- Missed spam
- Message marked as a spam, but is not a spam
- Missed marketing message
- Message marked as a marketing message, but is not a marketing message
- Missed phishing message

Related Topics

- [How to Report Incorrectly Classified Messages to Cisco](#), on page 25
- [How to Track Your Submissions](#), on page 30

How to Report Incorrectly Classified Messages to Cisco

Before You Begin

Before you start reporting incorrectly classified messages to Cisco, you must perform the following steps. Perform this step only once.

Procedure

Step 1

Register as an administrator on Cisco Talos Email Status Portal can be done in any one of the following ways:

Note Cisco Talos Email Status Portal is a web-based tool that allows email administrators to view and track email submissions on the portal.

- Registering when you are the first administrator in your organization to access the portal:
 - a. Log in to Cisco Talos Email Status Portal (https://talosintelligence.com/email_status_portal) using your Cisco credentials.
 - b. Click **Manage Account**.
 - c. Click **Add Domain**.
 - d. Enter your organization's domain name in the **Domain** field to register your domain with the portal.

Note Make sure that you enter a valid domain name, for example, `example.com` is the domain name in the following email address: `user@example.com`. If you have multiple domains in your organization, make sure that you add all the domains.

- e. Check the **I own this domain** check box if you are the owner of the domain entered in step 'd.'

Note If you do not check the 'I own this domain' check box, then you will only have domain view access rights. For more information, see the Cisco Talos Email Status Portal Help page at https://talosintelligence.com/tickets/email_submissions/help

- f. Click **Submit**.

After you click Submit, an email with a 6-digit character verification code is automatically sent to `postmaster@domain.com` (where `domain.com` is the domain you entered in step 'd') to confirm the domain ownership.

If your organization is not using `postmaster@domain.com` or your administrator does not have access to the postmaster mailbox, create a message filter (on all your appliances) to redirect messages from `SubmissionPortal@cisco.com` sent to `postmaster@domain.com` to a different email address. The following is a sample message filter:

```
redirect_postmaster: if (rcpt-to == "postmaster@domain.com") AND (mail-from ==
"^SubmissionPortal@cisco.com$") { alt-rcpt-to ("admin@domain.com"); }
```

- g. Enter the 6-digit character verification code in the **Domain Ownership Verification Code** dialog box to confirm the domain ownership.
- h. Click **Submit Verification Code**.

After you click the Submit Verification Code button, you are automatically provided with admin access rights. A registration ID is auto generated that can be viewed in the Manage Accounts section of the portal. You can use the registration ID for all the appliances in your organization.

Note A registration ID is a unique identifier to identify submissions made from the Cisco Email Security Gateways that belong to a particular organization.

- Registering when an administrator in your organization is already registered on the portal:

- a. Log in to Cisco Talos Email Status Portal (https://talosintelligence.com/email_status_portal) using your Cisco credentials.
- b. Click **Manage Account**.
- c. Click **Add Domain**.
- d. Enter your organization's domain name in the **Domain** field to register your domain with the portal.
Note Make sure that you enter a valid domain name, for example, `example.com` is the domain name in the following email address: `user@example.com`. If you have multiple domains in your organization, make sure that you add all the domains.
- e. Click **Submit**.

After you click Submit, an email notification is sent to the administrator who is already registered on the portal. This administrator must log in to the portal, and click **Approve** in the Permission Requests section of Manage Accounts to approve your registration request.

After your registration request is approved, a registration ID is auto generated that can be viewed in the Manage Accounts section of the portal. You can use the registration ID for all the appliances in your organization.

Note A registration ID is a unique identifier to identify submissions made from the Cisco Email Security Gateways that belong to a particular organization.

Step 2 Add the registration ID generated from Cisco Talos Email Status Portal for all the appliances in your organization.

- a. Log in to your appliances using the web interface.
- b. Go to **System Administration > Cisco Talos Email Status Portal Registration**.
- c. If your appliance is part of a cluster, set the mode to cluster level.
- d. Click **Set Registration ID**.
- e. Enter the registration ID obtained from the Cisco Talos Email Status Portal in the **Registration ID** field.
- f. Submit and commit your changes.
- g. If your appliance is not part of a cluster, you must repeat steps 1 through 6 on all the appliances in your organization.

You can also use the `portalregistrationconfig` command in CLI to set the registration ID.

How to Report Incorrectly Classified Messages to Cisco

For more information, see:

- How to Submit Email Messages to Cisco document at <https://www.cisco.com/c/en/us/support/docs/security/email-security-appliance/214133-how-to-submit-email-messages-to-cisco.html#anc5>.
- Cisco Talos Email Status Portal Help page at https://talosintelligence.com/tickets/email_submissions/help.

Procedure

Step 1 Perform the steps mentioned in **Before You Begin** section of [How to Report Incorrectly Classified Messages to Cisco](#), on page 25 .

Step 2 Report incorrectly classified messages to Cisco using one of the following methods:

- [Using Cisco Email Security Plug-In](#), on page 28
- [Forwarding Incorrectly Classified Message as an Attachment](#), on page 28

After you report an incorrectly classified message to Cisco, you will receive an email notification based on the option you select under the Email Notification and Reports button in the Manage Account section of the portal.

Note The 'My Submission Notifications' and 'My Submission Reports' options under the 'Email Notification and Reports' button are set to off by default. For more information, see the Cisco Talos Email Status Portal Help page at https://talosintelligence.com/tickets/email_submissions/help

What to do next

[How to Track Your Submissions](#), on page 30

Using Cisco Email Security Plug-In

Cisco Email Security Plug-In is a tool that allows users (email administrators and end users) to report incorrectly classified messages to Cisco using Microsoft Outlook. When you deploy this plug-in as part of Microsoft Outlook, a reporting menu is added to the Microsoft Outlook web interface. You can use the plug-in menu to report incorrectly classified messages.

Additional Information

- You can download Cisco Email Security Plug-In from the following page: <https://software.cisco.com/portal/pub/download/portal/select.html?&mdfid=284900944&flowid=41782&softwareid=283090986>.
- For more information, see the Cisco Email Security Plug-In Administrator Guide <http://www.cisco.com/c/en/us/support/security/email-encryption/products-user-guide-list.html>.

Forwarding Incorrectly Classified Message as an Attachment

Depending on the category of the message, you can forward each incorrectly classified message as an RFC 822 attachment to the following addresses as shown in the table below:

Email Submission	Definition	Submission Method	User Consideration for Submitting
Spam / Phish	Unsolicited and undesired. Spam/Phish is never legitimate and may also be malicious (phish, virus, malware, scams, etc.)	spam@access. ironport.com phish@access. ironport.com virus@access. ironport.com Outlook Plugin 'Spam', 'Phish', or 'Virus' button	Delivered to the user's inbox, but user considers message to be spam or Phish. Detected as spam, but user considers message legitimate.
Legitimate	Legitimate (good) email, not spam. Also known as 'Ham.'	ham@access. ironport.com Outlook Plugin 'Not Spam' button	Marketing/graymail messages not detected as marketing/graymail.
Marketing / Graymail	Marketing is legitimate (not Spam) email that is commercial bulk email. Usually subscription based, sometimes unwanted. Users may have knowingly or unknowingly solicited mail from the sender. For example swiping a badge at a conference or making an online purchase, etc. Legitimate subscription based marketing email will have a working unsubscribe mechanism. Graymail is a broader category that includes Marketing as well as other legitimate bulk mail.	ads@access. ironport.com Outlook Plugin 'Marketing' button	Detected as spam, but user considers message legitimate
Not Marketing / Graymail	Legitimate email (not Spam) that is not bulk and not subscription based. Usually person-to-person and/or transactional.	not_ads@access. ironport.com	Detected as Marketing/Graymail, but user considers the message to be transactional or otherwise not Marketing/Graymail.

You can achieve best results if you use one of the following email programs to forward the message:

- Apple Mail
- Microsoft Outlook for Mac
- Microsoft Outlook Web App
- Mozilla Thunderbird

**Caution**

If you are using Microsoft Outlook 2010, 2013, or 2016 for Microsoft Windows, you must use the Cisco Email Security Plug-In or the Microsoft Outlook Web App to report incorrectly classified messages. This is because Outlook for Windows may not forward the message with the required headers intact. Also, use the mobile platforms only if you can forward the original message as an attachment.

How to Track Your Submissions

After you receive an email notification with the submission details, you can view and track your submission on Cisco Talos Email Status Portal.

Procedure

- Step 1** Log in to Cisco Talos Email Status Portal (https://talosintelligence.com/email_status_portal) using your Cisco credentials.
- Step 2** Click **Submissions** on Cisco Talos Email Status Portal.
- Step 3** Click **Filter Options** and select appropriate filter option(s).
- Step 4** (Optional) Click the calendar button to choose the specific date.

What to do next

For more information, see the Cisco Talos Email Status Portal Help page at https://talosintelligence.com/tickets/email_submissions/help.

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 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 IP 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 appliance, as well as the header used to store the originating IP address.


Related Topics

- [Example Environments with Incoming Relays](#), on page 31
- [Configuring the Appliance to Work with Incoming Relays](#), on page 32
- [How Incoming Relays Affect Functionality](#), on page 37
- [Configuring Logs to Specify Which Headers Are Used](#), on page 39

Example Environments with Incoming Relays

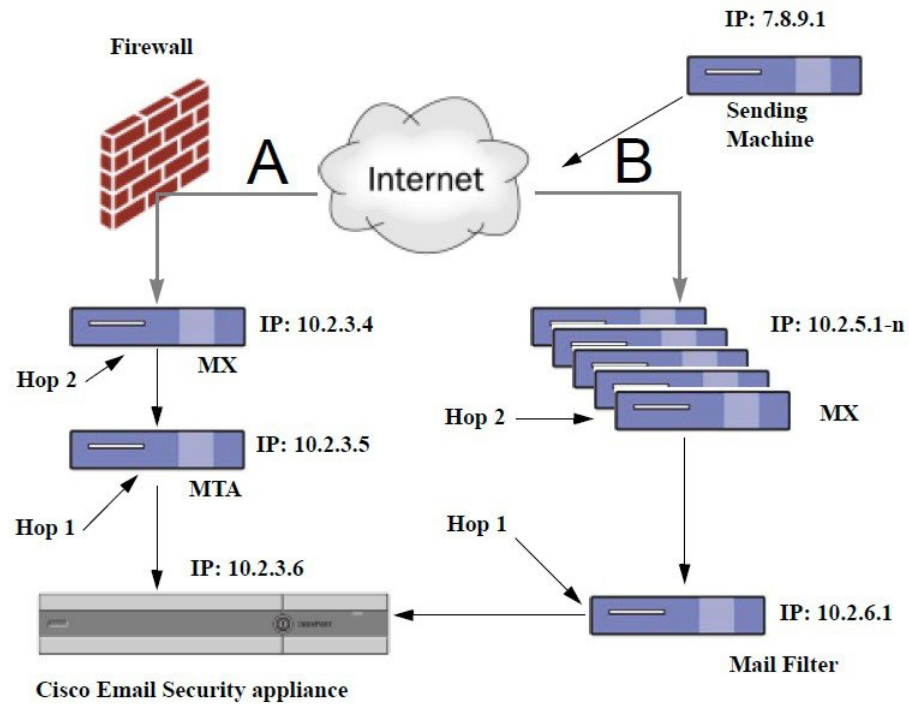
The following figure shows a very basic example of an incoming relay. Mail from IP address 7.8.9.1 appears to come from IP address 10.2.3.4 because the local MX/MTA is relaying mail to the appliance .

Figure 6: Mail Relayed by MX/MTA — Simple



The following figure 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 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 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 appliance .

Figure 7: Mail Relayed by MX/MTA — Advanced



Configuring the Appliance to Work with Incoming Relays

Related Topics

- [Enabling the Incoming Relays Feature](#) , on page 32
- [Adding an Incoming Relay](#) , on page 33
- [Message Headers for Relayed Messages](#) , on page 34

Enabling the Incoming Relays Feature



Note You should only enable the incoming relays feature if a local MX/MTA relays mail to your 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 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](#) , on page 34.

- 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 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, on page 38](#).
- To facilitate tracking and troubleshooting, configure the appliance logs to show which header is used. See [Configuring Logs to Specify Which Headers Are Used , on page 39](#).

Related Topics

- [How to Configure the Appliance to Scan Messages for Spam, on page 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 , on page 34](#)
- [Received Header, on page 35](#)

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 the following figure, 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 8: Mail Relayed by MX/MTA — Variable Number of Hops



Related Topics

- [Adding an Incoming Relay](#) , on page 33

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 - Mail Relayed by MX/MTA — Advanced*) 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 - Mail Relayed by MX/MTA — Variable Number of Hops*) to the machine connecting to your appliance , you must use a custom header (see [Custom Header](#) , on page 34).

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 appliance does not count as a hop. See [Configuring Logs to Specify Which Headers Are Used](#) , on page 39 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 appliance is parsed. If neither the parsing character nor a valid IP address is found, the 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 - Mail Relayed by MX/MTA — Advanced* 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)

The following table shows example email headers for a message as it moves through several hops on its way to the appliance as in *Figure - Mail Relayed by MX/MTA — Advanced*. This example shows extraneous headers (ignored by your appliance) which are present once the message has arrived in the recipient's inbox. The number of hops to specify would be two.

Table 1: A Series of Received: Headers (Path A Example 1)

1	<pre>Microsoft Mail Internet Headers Version 2.0 Received: from smemail.rand.org ([10.2.2.7]) by smmail5.customerdoamin.org with Microsoft SMTPSVC(5.0.2195.6713); Received: from ironport.customerdomain.org ([10.2.3.6]) by smemail.customerdoamin.org with Microsoft SMTPSVC(5.0.2195.6713);</pre>
2	<pre>Received: from mta.customerdomain.org ([10.2.3.5]) by ironport.customerdomain.org with ESMTP; 21 Sep 2005 13:46:07 -0700</pre>
3	<pre>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></pre>
4	<pre>Received: from sending-machine.spamham.com (sending-machine.spamham.com [7.8.9.1]) by mx.customerdomain.org (Postfix) with ESMTP id 4F3DA15AC22 for <joefoo@customerdomain.org></pre>
5	<pre>Received: from linux1.thespammer.com (HELO linux1.thespammer.com) ([10.1.1.89]) by sending-machine.spamham.com with ESMTP; Received: from exchange1.thespammer.com ([10.1.1.111]) by linux1.thespammer.com with Microsoft SMTPSVC(6.0.3790.1830); Subject: Would like a bigger paycheck? Date: Wed, 21 Sep 2005 13:46:07 -0700 From: "A. Sender" <asend@otherdomain.com> To: <joefoo@customerdomain.org></pre>

Notes for the above table:

- The appliance ignores these headers.
- The 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 appliance ignores these Microsoft Exchange headers.

The following table shows the headers for the same email message, without the extraneous headers

Table 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>;

The following figure shows the incoming relay for path A (above) as configured in the Add Relay page in the GUI:

Figure 9: A Configured Incoming Relay with Received Header

Add Relay

The screenshot shows the 'Add Relay' configuration interface. Under the 'Incoming Relay' section, the following fields are visible:

- Name:** IncomingRelayOne
- IP Address:** 10.2.3.5
- Header:** Two radio buttons are present. The first is 'Specify a custom header' (unselected). The second is 'Parse the "Received" header' (selected).
- Begin parsing after:** A text input field containing '['.
- Hop:** A dropdown menu set to '2'.

Related Topics

- [Adding an Incoming Relay](#) , on page 33

How Incoming Relays Affect Functionality

- [Incoming Relays and Filters](#), on page 37
- [Incoming Relays, HAT, IP Reputation Score, and Sender Groups](#), on page 38
- [Incoming Relays and Directory Harvest Attack Prevention](#), on page 38
- [Incoming Relays and Trace](#), on page 38
- [Incoming Relays and Email Security Monitor \(Reporting\)](#) , on page 38
- [Incoming Relays and Message Tracking](#), on page 38
- [Incoming Relays and Logging](#) , on page 38

Incoming Relays and Filters

The Incoming Relays feature provides the various IP Reputation Service related filter rules (reputation, no-reputation) with the correct IP Reputation score.

Incoming Relays, HAT, IP Reputation Score, and Sender Groups

HAT policy groups do not currently use information from Incoming Relays. However, because the Incoming Relays feature does supply the 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 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 IP 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 IP 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 IP 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 IP Reputation score for the sender is reported initially on line 1. Later, once the Incoming Relay is processed, the correct IP Reputation score is reported on line 5.

1	Fri Apr 28 17:07:29 2006 Info: ICID 210158 ACCEPT SG UNKNOWNLIST match nx.domain IPR 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, IPR 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 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 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](#).

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 > IMS and Graymail**.
- 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 .

What to do next

Related Topics

- [Service Updates](#)
- [Updates Through a Proxy Server](#)
- [Configuring Server Settings for Downloading Upgrades and Updates](#)

Testing Anti-Spam

To	Do This	More Information
Test your configuration.	<p>Test your configuration using the <code>X-advertisement: spam</code> header.</p> <p>For testing purposes, Cisco Anti-Spam considers any message with an X-header formatted as <code>X-Advertisement: spam to be spam</code>.</p>	<p>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, on page 18) are performed.</p> <p>Use this header with one of the following:</p> <ul style="list-style-type: none"> • Use SMTP commands to send a test message with this header. See Sending an Email to the Appliance to Test Cisco Anti-Spam, on page 41. • Use the trace command and include this header. See Debugging Mail Flow Using Test Messages: Trace.
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 , on page 42.

Related Topics

- [Sending an Email to the Appliance to Test Cisco Anti-Spam](#), on page 41
- [Ways Not to Test Anti-Spam Efficacy](#), on page 42

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

Before You Begin

Review the example in [Testing Anti-Spam Configuration: Example Using SMTP, on page 41](#).

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, on page 41](#)
-

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 ESMTTP
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 IP Reputation Service, blocked lists, 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.