



Understanding the Email Pipeline

This chapter contains the following sections:

- [Overview of the Email Pipeline, on page 1](#)
- [Email Pipeline Flows, on page 1](#)
- [Incoming / Receiving, on page 4](#)
- [Work Queue / Routing, on page 6](#)
- [Delivery, on page 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

The following figures provide an overview of how email is processed through the system, from receipt 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 1: Email Pipeline — Receiving Email Connections




Figure 2: Email Pipeline — Work Queue





Figure 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, on page 4](#)
- [Received: Header, on page 5](#)
- [Default Domain, on page 5](#)
- [Bounce Verification, on page 5](#)
- [Domain Map, on page 5](#)
- [Recipient Access Table \(RAT\), on page 5](#)
- [Alias Tables, on page 5](#)
- [LDAP Recipient Acceptance, on page 6](#)
- [SMTP Call-Ahead Recipient Validation, on page 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 SenderBase Reputation Service.

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

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 [Working with Listeners](#).

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 [Working with Listeners](#).

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](#).

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](#).

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](#).

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](#).

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. For more information, see [Working with Listeners](#). 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 [Working with LDAP Queries](#).

SMTP Call-Ahead Recipient Validation

When you configure your appliance for SMTP call-ahead recipient validation, the 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 appliance. The 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 [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](#).

Related Topics

- [Email Pipeline and Security Services, on page 6](#)
- [LDAP Recipient Acceptance, on page 6](#)
- [Masquerading or LDAP Masquerading, on page 7](#)
- [LDAP Routing, on page 7](#)
- [Message Filters, on page 7](#)
- [Email Security Manager \(Per-Recipient Scanning\), on page 8](#)
- [Quarantines, on page 9](#)

Email Pipeline and Security Services

It is recommended that you enable and avoid changing security services on Cloud Email Security appliances

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. For more information, see [Working with Listeners](#). 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 [Working with LDAP Queries](#).

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](#).

For more information about masquerading via an LDAP query, see [Working with LDAP Queries](#).

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 [Working with LDAP Queries](#).

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 [Using Message Filters to Enforce Email Policies](#).

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, on page 8](#)
- [Anti-Spam, on page 8](#)
- [Anti-Virus, on page 8](#)
- [Graymail Detection and Safe Unsubscribing, on page 9](#)
- [File Reputation Scanning and File Analysis , on page 9](#)
- [Content Filters, on page 9](#)
- [Outbreak Filters, on page 9](#)

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 [Spam Quarantine](#).

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.

For more information, see [Managing Spam and Graymail](#) .

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, on page 9](#)) are scanned for viruses. For more information about Anti-Virus scanning, see [Anti-Virus](#).

Graymail Detection and Safe Unsubscribing

You can configure the appliance to detect graymail messages and perform secure unsubscribe on behalf of the end user. Available actions are similar to those for anti-virus scanning.

For more information, see [Managing Spam and Graymail](#).

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 [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 [Content Filters](#).

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, on page 9](#).

For more information, see [Outbreak Filters](#).

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

- [Policy, Virus, and Outbreak Quarantines](#)
- [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, on page 10](#)
- [Delivery Limits, on page 10](#)
- [Domain-Based Limits, on page 10](#)
- [Domain-Based Routing, on page 10](#)
- [Global Unsubscribe, on page 11](#)
- [Bounce Limits, on page 11](#)

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 [Configuring Mail Gateways for all Hosted Domains Using Virtual Gateway™ Technology](#).

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](#).

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 Using Destination Controls](#).

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](#).

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](#).

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](#).

